

Monitoring a node during a job.

The Delta documentation has moved to <https://docs.ncsa.illinois.edu/systems/delta/>. Please update any bookmarks you may have. Click in the link above if you are not automatically redirected in 5 seconds.

You have ssh access to nodes in your running job(s) . A couple basic monitoring tools are demonstrated in the example transcript here. Screen shots are appended so that you can see the output from the tools. Most common linux utilities are available from the compute nodes (free, strace, ps ...).

```
[arnoldg@dt-login03 python]$ squeue -u $USER
      JOBID PARTITION   NAME     USER ST       TIME  NODES NODELIST(REASON)
      1214412 gpuA40x4- interact arnoldg R        8:14    1 gpub045
[arnoldg@dt-login03 python]$ ssh gpub045
gpub045.delta.internal.ncsa.edu (141.142.145.145)
  OS: RedHat 8.4  HW: HPE  CPU: 64x  RAM: 252 GB
Last login: Wed Dec 14 09:45:26 2022 from 141.142.144.42
[arnoldg@gpub045 ~]$ nvidia-smi

[arnoldg@gpub045 ~]$ module load nvtop
-----
The following dependent module(s) are not currently loaded: cuda/11.6.1 (required by: ucx/1.11.2, openmpi/4.1.2)
-----

The following have been reloaded with a version change:
  1) cuda/11.6.1 => cuda/11.7.0

[arnoldg@gpub045 ~]$ nvtop

[arnoldg@gpub045 ~]$ module load anaconda3_gpu
[arnoldg@gpub045 ~]$ nvtop

[arnoldg@gpub045 ~]$ top -u $USER
```

nvidia-smi :

jupyter.salloc - python [SSH: delta] - Visual Studio Code

File Edit Selection View Go Run Terminal Help

jobcharge_grp.py 1, 4 jupyter.salloc

```

1 salloc --account=bbka-delta-gpu --partition=gpuA40x4-interactive \
2   --nodes=1 --tasks-per-node=1 \
3   --cpus-per-task=2 --mem=16g \
4   --gpus-per-node=2
5
6

```

PROBLEMS 1 OUTPUT DEBUG CONSOLE TERMINAL PORTS

```

+-----+
[arnoldg@gpub045 ~]$ nvidia-smi
Wed Dec 14 09:46:46 2022
+-----+
| NVIDIA-SMI 515.65.01    Driver Version: 515.65.01    CUDA Version: 11.7     |
+-----+
| GPU  Name            Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
|=====-=|
|  0  NVIDIA A40       On          | 00000000:46:00:0 Off  |    24%    Default  |
|    34C   P0       79W / 300W | 259MiB / 46068MiB |             N/A     |
+-----+
|  1  NVIDIA A40       On          | 00000000:C7:00:0 Off  |    0%     Default  |
|    33C   P0       79W / 300W | 259MiB / 46068MiB |             N/A     |
+-----+
+-----+
| Processes: |
| GPU  GI  CI           PID  Type  Process name                        GPU Memory |
|=====-=|
|  0  N/A N/A         2665497  C    ...x/release/./bandwidthTest       257MiB |
|  1  N/A N/A         2665497  C    ...x/release/./bandwidthTest       257MiB |
+-----+
[arnoldg@gpub045 ~]$

```

SSH: delta master* 0 1 0

Ln 4, Col 20 Spaces: 2 UTF-8 LF Plain Text

nvtop :

jupyter.salloc - python [SSH: delta] - Visual Studio Code

File Edit Selection View Go Run Terminal Help

jobcharge_grp.py 1, 4 jupyter.salloc

```

1 salloc --account=bbka-delta-gpu --partition=gpuA40x4-interactive \
2   --nodes=1 --tasks-per-node=1 \
3   --cpus-per-task=2 --mem=16g \
4   --gpus-per-node=2
5
6

```

PROBLEMS 1 OUTPUT DEBUG CONSOLE TERMINAL PORTS

Device 0 [NVIDIA A40] PCIe GEN 4@16x RX: 0.000 KiB/s TX: 0.000 KiB/s
GPU 210MHz MEM 7250MHz TEMP 33°C FAN 0% POW 57 / 300 W
GPU[0%] MEM[0.560Gi/44.988Gi]

Device 1 [NVIDIA A40] PCIe GEN 4@16x RX: 0.000 KiB/s TX: 0.000 KiB/s
GPU 210MHz MEM 7250MHz TEMP 33°C FAN 0% POW 57 / 300 W
GPU[0%] MEM[0.560Gi/44.988Gi]

PID	USER	DEV	TYPE	GPU	GPU MEM	CPU	HOST MEM	Command
F2	Setup							
F6	Sort							
F9	Kill							
F10	Quit							
F12	Save Config							

F2 Setup F6 Sort F9 Kill F10 Quit F12 Save Config

SSH: delta master* 0 0 1 0 Ln 4, Col 20 Spaces: 2 UTF-8 LF Plain Text

nvitop :

jupyter.salloc - python [SSH: delta] - Visual Studio Code

File Edit Selection View Go Run Terminal Help

jobcharge_grp.py 1, 4 jupyter.salloc

```

1 salloc --account=bbka-delta-gpu --partition=gpuA40x4-interactive \
2 --nodes=1 --tasks-per-node=1 \
3 --cpus-per-task=2 --mem=16g \
4 --gpus-per-node=2
5
6

```

PROBLEMS 1 OUTPUT DEBUG CONSOLE TERMINAL PORTS

Wed Dec 14 09:50:59 2022 (Press *h* for help or *q* to quit)

NVITOP 0.10.2		Driver Version: 515.65.01		CUDA Driver Version: 11.7	
GPU	Name	Persistence-M	Bus-Id	Disp.A	Volatile Uncorr. ECC
Fan	Temp	Perf	Pwr:Usage/Cap	Memory-Usage	GPU-Util Compute M.
0	A40	On	00000000:46:00:0	Off	0
0%	35C	P0	326MiB / 44.99GiB		26%
MEM: 0.7%		UTL: 26%			
1	A40	On	00000000:C7:00:0	Off	0
0%	33C	P8	3MiB / 44.99GiB		0%
MEM: 0.0%		UTL: 0%			

Load Average: 3.10 2.92 2.85
CPU: 5.7%

MEM: 13.2%
SWP: 0.0%

AVG GPU MEM: 0.3%

AVG GPU UTL: 28.5%

Processes:									
GPU	PID	USER	GPU-MEM	%SM	%CPU	%MEM	TIME	COMMAND	
0	2666414	C arnoldg	311MiB	53	99.8	0.1	0:06	/projects/bbka/slurm_test_scripts/gpu/cuda/cuda..	arnoldg@gpub045.delta.internal.ncsa.edu

SSH: delta master* 0 0 1 0 Ln 4, Col 20 Spaces: 2 UTF-8 LF Plain Text

top -u \$USER :

jobcharge_grp.py 1, 4

jupyter.salloc

jupyter.salloc

```

1  salloc --account=bbka-delta-gpu --partition=gpuA40x4-interactive \
2  --nodes=1 --tasks-per-node=1 \
3  --cpus-per-task=2 --mem=16g \
4  --gpus-per-node=2
5
6
    
```

PROBLEMS 1 OUTPUT DEBUG CONSOLE TERMINAL PORTS

```

top - 09:54:25 up 120 days, 20:33, 6 users, load average: 3.17, 2.94, 2.86
Tasks: 999 total, 4 running, 995 sleeping, 0 stopped, 0 zombie
%Cpu(s): 5.0 us, 0.9 sy, 0.0 ni, 94.1 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
MiB Mem : 257667.4 total, 201574.9 free, 19102.0 used, 36990.5 buff/cache
MiB Swap: 0.0 total, 0.0 free, 0.0 used. 223413.5 avail Mem
    
```

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
2666701	arnoldg	20	0	9982.2m	343032	288896	R	100.0	0.1	0:23.85	bandwidthTest
2666709	arnoldg	20	0	276164	6132	4420	R	0.3	0.0	0:00.03	top
2666452	arnoldg	20	0	177476	6692	4720	S	0.0	0.0	0:00.00	sshd
2666453	arnoldg	20	0	237576	7760	3520	S	0.0	0.0	0:00.02	bash

ssh
srun release