

# SambaNova DataScale SN30: Getting Started

**Varuni Sastry**  
Data Science and Learning  
[vsastry@anl.gov](mailto:vsastry@anl.gov)

July 11, 2023

# OVERVIEW

- System Overview
- Account access
- Steps to Login
- Environment setup
- Workflow
- Example programs


# SambaNova Datascale SN30

<https://www.alcf.anl.gov/alcf-ai-testbed>



SambaNova Datascale SN30

- 4 Racks
- 8 nodes of SN30
- 8 RDUs or 4 XRDU's per node
- 8 Tiles per RDU
- Group of 4 tiles



# SambaNova Datascale SN30

<https://www.alcf.anl.gov/alcf-ai-testbed>



SambaNova Datascale SN30

- 4 Racks
- 8 nodes of SN30
- 8 RDUs or 4 XRDU's per node
- 8 Tiles per RDU
- Group of 4 tiles



# ACCOUNT ACCESS

- **Step 1 : Request for an ALCF account**

Support Center | ALCF Resources | ALCF Website

Argonne NATIONAL LABORATORY | Argonne Leadership Computing Facility

Home

Request an Account

Reactivate an Account

Allocation Requests

ALCF Passcode Token Help

## ALCF Account and Project Management

### ALCF USER SUPPORT

ACCOUNTS RELATED HELP: [accounts@alcf.anl.gov](mailto:accounts@alcf.anl.gov)  
SUPPORT TICKETS: [support@alcf.anl.gov](mailto:support@alcf.anl.gov)  
TELEPHONE: Email [accounts@alcf.anl.gov](mailto:accounts@alcf.anl.gov) for the phone support number

ALCF HELP DESK ADDRESS:  
Argonne Leadership Computing Facility  
9700 S. Cass Avenue  
Building 240, #2129  
Lemont, IL 60439

#### Account Log In

Username [Forgot username?](#)

CRYPTOCard Passcode [Token issues?](#)

**LOGIN**

[Request a new account](#)

Request an ALCF account on our [accounts page](#).

Need an active project account.

Help : [accounts@alcf.anl.gov](mailto:accounts@alcf.anl.gov)

# ACCOUNT ACCESS

- **Step 2 : Request to Join project**

Home

Accounts ▾

Projects & Resources ▾

Join project

Request and view systems

Manage UNIX Groups

Request an Allocation

## Join project

Request to become a team member of an existing project.

### Available Projects

Click on the name of the project to add and remove proxies and team members.

sn_	Filter by Title	Filter by PI
Project Name ▾	Title ▾	PI ▾
sn_training	SambaNova Training	Venkatram Vishwanath

1 total

Join project under "sn\_training".

## Request and view systems

Request to use additional systems.

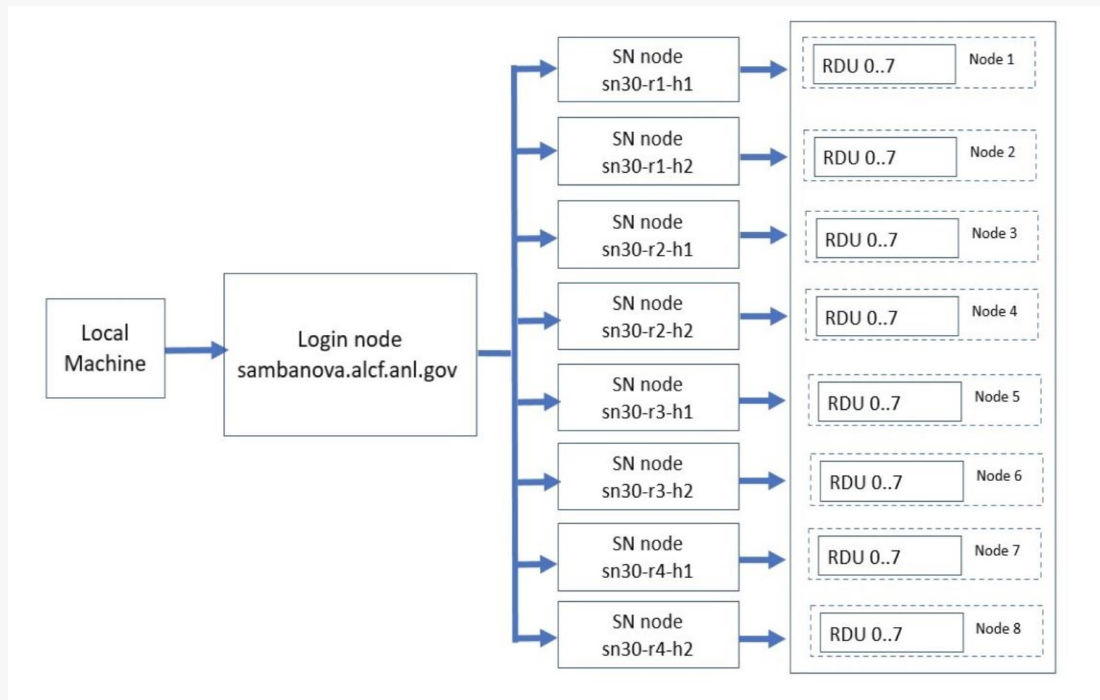
### Available Resources

Current: sambanova.

Check for "sambanova" under Request and view systems


# Connection and Login

- **Step 3 : Login to SN30 node.**



From local machine to login node :

```
ssh ALCFUserID@sambanova.alcf.anl.gov  
Password: < MobilePASS+ code >
```



From login node to a SN30 node :

```
ALCFUserId@sm-login-1:~$ ssh sn30-r1-h1
```

# Environment setup

- SDK (SambaFlow software stack and the associated environmental variables) is set up at login.
- Pre-built environments under `/opt/sambaflow/apps/`
- Create new virtual environment and install packages

```
python -m venv --system-site-packages my_env  
source my_env/bin/activate
```

```
python3 -m pip install <package>
```

```
./private/an1/venv  
./recommender/dlrm/venv  
./recommender/deepinterest/venv  
./recommender/ncf/venv  
./starters/lenet/venv  
./starters/upscalenet/venv  
./starters/mlp/venv  
./starters/ffn_mnist/venv  
./starters/power_pca/venv  
./starters/logreg/venv  
./micros/venv  
./nlp/tapas/venv  
./nlp/data_processing/venv  
./nlp/transformers_on_rdu/blocksparse/venv  
./nlp/transformers_on_rdu/genslm/venv  
./nlp/transformers_on_rdu/gpt13b/venv  
./nlp/transformers_on_rdu/venv  
./speech/speaker_diarization/venv  
./speech/hubert/venv  
./image/deepvit/venv  
./image/object_detection/venv  
./image/segmentation_3d/venv  
./image/segmentation/venv  
./image/classification/venv
```



# Workflow

- **Compile**

- Model compilation and '.pef' generation.
- Maps the compute and memory resources required to run an application on RDUs
- Re-compile only when model parameters change.
- Significant compile times for large models.

```
srun python lenet.py compile -b=1 --pef-name="lenet" --output-folder="pef"
```

- **Run**

- Model trained on RDU using the ".pef" generated as part of compile process and the training dataset.

```
srun python lenet.py run --pef="pef/lenet/lenet.pef"
```

# Example programs

- /opt/sambaflow/apps/

```
drwxr-xr-x 7 root root 4096 Mar 27 04:15 image → drwxr-xr-x 8 root root 4096 Mar 27 04:15 deepvit
drwxr-xr-x 5 root root 4096 Mar 27 04:16 nlp          drwxr-xr-x 5 root root 4096 Mar 27 04:15 object_detection
drwxr-xr-x 6 root root 4096 Mar 27 04:17 micros      drwxr-xr-x 5 root root 4096 Mar 27 04:15 classification
drwxr-xr-x 5 root root 4096 Mar 27 04:17 recommender drwxr-xr-x 8 root root 4096 Mar 27 04:16 segmentation_3d
drwxr-xr-x 4 root root 4096 Mar 27 04:17 speech     drwxr-xr-x 5 root root 4096 Mar 27 21:09 segmentation
drwxr-xr-x 8 root root 4096 Mar 27 04:18 starters
```

- /data/ANL/scripts

```
-rwxrwxr-x 1 root SambaNova_Deployment 4284 Mar 31 18:13 Gpt1.5B_compile.sh
-rwxrwxr-x 1 nobody SambaNova_Deployment 4391 Mar 31 18:13 Gpt1.5B_compile_single.sh
-rwxrwxr-x 1 root SambaNova_Deployment 4491 Mar 31 18:13 Gpt1.5B_run.sh
-rwxrwxr-x 1 root SambaNova_Deployment 4658 Mar 31 18:13 Gpt1.5B.sh
-rwxrwxr-x 1 nobody SambaNova_Deployment 2160 Mar 31 18:14 Stream-DDR-BW.sh
-rwxrwxr-x 1 root SambaNova_Deployment 2183 Mar 31 18:14 Stream-Host-BW.sh
-rwxrwxr-x 1 root SambaNova_Deployment 2732 Mar 31 18:14 Unet3D-512.sh
-rwxrwxr-x 1 root SambaNova_Deployment 2935 Mar 31 18:14 Unet3d.sh
-rwxrwxr-x 1 root SambaNova_Deployment 2283 Mar 31 18:14 Uno.sh
```

# Example program - mnist

- Path : "/opt/sambaflow/apps/starters/ffn\_mnist"  
**\*Note : Make a copy of the apps directory into the home directory.**
- Activate the environment.

```
source /opt/sambaflow/apps/starters/ffn_mnist/venv/bin/activate
```

- Compile and Run

```
srun python ffn_mnist.py compile -b 1 --pef-name="ffn_mnist" --mac-v2
```

```
srun python ffn_mnist.py run -b 1 -p out/ffn_mnist/ffn_mnist.pef
```

# Example program – multinode

- Runs multiple instances of training on multiple tiles / RDUs / nodes in a data-parallel fashion.
- Gradient accumulation on the RDU.
- Refer example :
  - /data/ANL/scripts/Unet2d.sh
  - /data/ANL/scripts/unet\_batch.sh
  - [https://docs.alcf.anl.gov/ai-testbed/sambanova\\_gen2/example-multi-node-programs/](https://docs.alcf.anl.gov/ai-testbed/sambanova_gen2/example-multi-node-programs/)
- Environmental Variables
  - OMP\_NUM\_THREADS (8/16/32)
  - SF\_RNT\_NUMA\_BIND (2)
  - SF\_RNT\_FSM\_POLL\_BUSY\_WAIT (1)
  - SF\_RNT\_DMA\_POLL\_BUSY\_WAIT (1)
  - SAMBA\_CCL\_USE\_PCIE\_TRANSPORT (1)

# Example program – multinode

For single instance run :

```
./Unet2d.sh compile <image size> <batch_size> <num of instances> <RunID>  
./Unet2d.sh run <image size> <batch_size> <num of instances> <RunID>
```

```
./Unet2d.sh compile 256 256 1 unet2d_single_compile  
./Unet2d.sh run 256 256 1 unet2d_single_run
```

For multi instance run :

```
./Unet2d.sh pcompile <image size> <batch_size> <num of instances> <RunID>  
./Unet2d.sh prun <image size> <batch_size> <num of instances> <RunID>
```

```
./Unet2d.sh pcompile 256 256 8 unet2d_8inst_pcompile  
./Unet2d.sh prun 256 256 8 unet2d_8inst_prun
```

Note : Run the scripts after copying to your home directory

# Example program – multinode

- Activate environment

```
source /opt/sambaflow/apps/image/segmentation/venv/bin/activate
```

- Compile

```
python /opt/sambaflow/apps/image/segmentation/compile.py compile --mac-v2 --in-channels=3 --in-width=${2} --in-height=${2} --batch-size=${BS} --enable-conv-tiling --num-tiles=4 --pef-name=unet_train_${BS}_${2}_NP_${NUM_TILES} --data-parallel -ws 2 --output-folder=${OUTDIR}
```

- Run

```
sbatch --gres=rdu:1 --tasks-per-node ${NP} --nodes 1 --odelist $(hostname) --cpus-per-task=${cpus} $(pwd)/unet_batch.sh ${NP} ${NUM_WORKERS} ${BS} ${2} ${5}
```

```
srun --mpi=pmi2 python /opt/sambaflow/apps/image/segmentation//hook.py run --data-cache=${CACHE_DIR} --data-in-memory --num-workers=${NUM_WORKERS} --enable-tiling --min-throughput 395 --in-channels=3 --in-width=${IM} --in-height=${IM} --init-features 32 --batch-size=${BS} --epochs 10 --data-dir ${DS} --log-dir log_dir_unet_${IM}_${BS}_${NP} --data-parallel --reduce-on-rdu --pef=${OUTDIR}/unet_train_${BS}_${IM}_NP_4/unet_train_${BS}_${IM}_NP_4.pef
```

# Utility commands

**srun / sbatch** : Run individual Python scripts in parallel with other scripts on cluster assigned by Slurm.

```
srun --odelist=sn30-r1-h1 python lenet.py compile -b=1 --pef-name="lenet" --output-folder="pef"
```

**squeue** : Check the scheduled jobs

```
@sn30-r1-h1:~$ squeue
JOBID PARTITION NAME USER ST TIME NODES NODELIST(REASON)
14340 samanova python ████████ R 0:06 1 sn30-r1-h1
```

**sntilestat** : Check the process on each tile / node.

```
████████@sn30-r1-h1:~$ sntilestat
TILE %idle %exec %pload %aload %chkpt %quiesce PID USER COMMAND
/XRDU_0/RDU_0/TILE_0 82.7 3.7 10.7 2.9 0.0 0.0 492994 /opt/sambaflow/apps/starters/ffn_mnist/venv/bin/py
/XRDU_0/RDU_0/TILE_1 83.0 2.9 11.2 3.0 0.0 0.0 492994 /opt/sambaflow/apps/starters/ffn_mnist/venv/bin/py
/XRDU_0/RDU_0/TILE_2 83.9 2.7 10.9 2.4 0.0 0.0 492994 /opt/sambaflow/apps/starters/ffn_mnist/venv/bin/py
/XRDU_0/RDU_0/TILE_3 82.3 3.5 11.3 2.9 0.0 0.0 492994 /opt/sambaflow/apps/starters/ffn_mnist/venv/bin/py
/XRDU_0/RDU_0/TILE_4 83.9 2.4 10.9 2.7 0.0 0.0 492994 /opt/sambaflow/apps/starters/ffn_mnist/venv/bin/py
/XRDU_0/RDU_0/TILE_5 84.2 1.5 11.2 3.1 0.0 0.0 492994 /opt/sambaflow/apps/starters/ffn_mnist/venv/bin/py
/XRDU_0/RDU_0/TILE_6 84.1 1.1 11.7 3.1 0.0 0.0 492994 /opt/sambaflow/apps/starters/ffn_mnist/venv/bin/py
/XRDU_0/RDU_0/TILE_7 83.8 2.0 12.0 2.2 0.0 0.0 492994 /opt/sambaflow/apps/starters/ffn_mnist/venv/bin/py
/XRDU_0/RDU_1/TILE_0 100.0 0.0 0.0 0.0 0.0 0.0
/XRDU_0/RDU_1/TILE_1 100.0 0.0 0.0 0.0 0.0 0.0
```

**sinfo , scancel**

# Important directory paths and links

- Sambaflow models path : /opt/sambaflow/apps/
- Model scripts : /data/ANL/scripts/
- Important datasets : /software/sambanova/dataset/
- /software /projects
- [AI Testbed User Guide](#)
- [Sambanova Documentation.](#)
- [SambaTune.](#)



# Allocation programs



Director's Discretionary (DD) awards support various project objectives from scaling code to preparing for future computing competition to production scientific computing in support of strategic partnerships.

SN30 system is available for Director's Discretionary (DD) allocations.

[Allocation Request Form](#)

Allocations also available under [INCITE](#), [ALCC](#), and [ALCF Data Science](#) programs

THANK YOU

[www.anl.gov](http://www.anl.gov)

# BACKUP

[www.anl.gov](http://www.anl.gov)

# ACCOUNT ACCESS

- **Step 2 : Request for an ALCF account**

Argonne National Laboratory | Argonne Leadership Computing Facility

Support Center | ALCF Resources | ALCF Website

## ALCF Account and Project Management

### ALCF USER SUPPORT

ACCOUNTS RELATED HELP: [accounts@alcf.anl.gov](mailto:accounts@alcf.anl.gov)  
SUPPORT TICKETS: [support@alcf.anl.gov](mailto:support@alcf.anl.gov)  
TELEPHONE: Email [accounts@alcf.anl.gov](mailto:accounts@alcf.anl.gov) for the phone support number

ALCF HELP DESK ADDRESS:  
**Argonne Leadership Computing Facility**  
9700 S. Cass Avenue  
Building 240, #2129  
Lemont, IL 60439

#### Account Log In

Username [Forgot username?](#)

CRYPTOCARD Passcode [Token issues?](#)

**LOGIN**

[Request a new account](#)

Request an ALCF account on our [accounts page](#).

Need an active project account.

Help : [accounts@alcf.anl.gov](mailto:accounts@alcf.anl.gov)

# ACCOUNT ACCESS

- **Step 2 : Request for an ALCF account**

The screenshot shows the 'Join project' page. On the left is a blue navigation sidebar with links: Home, Accounts, Projects & Resources, Join project, Request and view systems, Manage UNIX Groups, and Request an Allocation. The main content area is titled 'Join project' and includes the instruction 'Request to become a team member of an existing project.' Below this is the 'Available Projects' section, which contains a search bar with 'sn\_' entered, and filter boxes for 'Filter by Title' and 'Filter by PI'. A table lists available projects:

Project Name	Title	PI
<a href="#">sn_training</a>	SambaNova Training	Venkatram Vishwanath

At the bottom of the table, it says '1 total'.

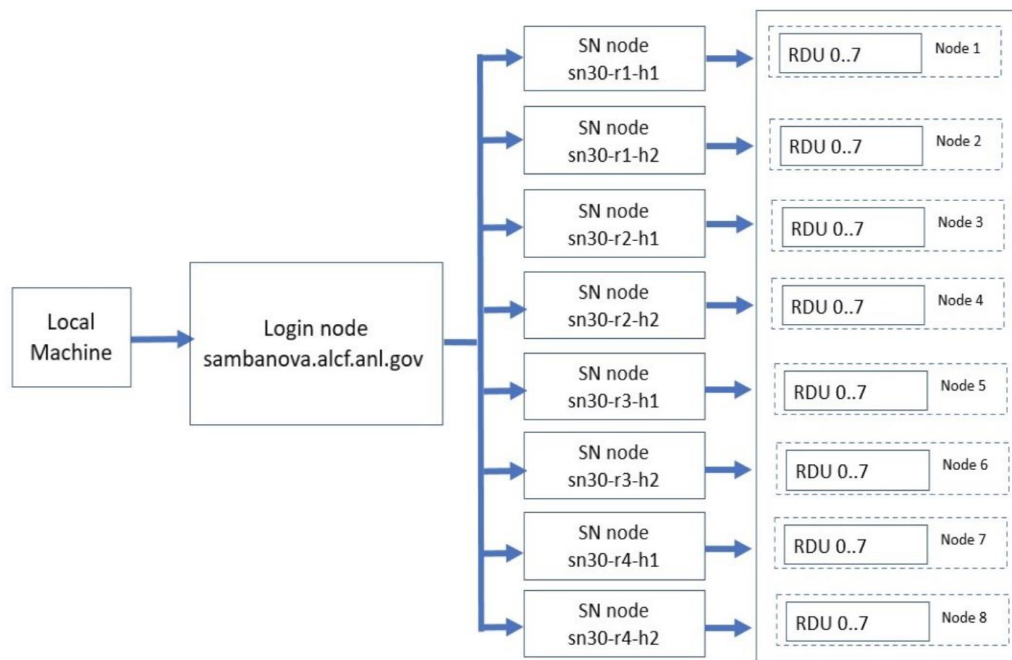
Join project under "sn\_training".

The screenshot shows the 'Request and view systems' page with the instruction 'Request to use additional systems.' Below is the 'Available Resources' section. A horizontal bar shows the current resource 'sambanova,' which is circled in red.

Check for "sambanova" under Request and view systems


# Connection and Login

- **Step 3 : Login to SN30 node.**



From local machine to login node :

```
ssh ALCFUserID@sambanova.alcf.anl.gov  
Password: < MobilePASS+ code >
```



From login node to a SN30 node :

```
ALCFUserId@sm-login-1:~$ ssh sn30-r1-h1
```