

Azure Batch AI – Walkthrough and How it works

BY [TSUYOSHI MATSUZAKI](#) ON [2018-01-30](#) • ([4 COMMENTS](#))

In my [previous post](#), I showed how to setup and configure for the training with multiple machines. Using Azure Batch AI, you'll find that it's very straightforward and simplifying many tasks for AI workloads, especially when it's distributed training workloads.

Azure Batch AI is built on top of Azure Batch (see my [early post](#)), and then you can easily understand this mechanism when you're familiar with Azure Batch.

Now let's see how Batch AI works and how you can use it !

Here we use Azure CLI (Command Line Interface), but it's also integrated with tool's UI like Azure Portal, Visual Studio, or Visual Studio Code.

Note : You can also use Azure Databricks (which is fully-managed service built on top of Spark) for distributed deep learning workloads in Azure.

Simplifies the distributed training

First, we assume that you are familiar with the basic concepts of distributed training with Cognitive Toolkit (CNTK). In this post we use the following distributed-training sample code (Train_MNIST.py) with training data (Train-28x28_cntk_text.txt), which is the exactly same sample code in my previous post. (See "[Walkthrough – Distributed training with CNTK](#)" for details.)

This sample trains the convolutional neural network (LeNet) with distributed manners and save the trained model (ConvNet_MNIST.dnn) after the training is done.

Train_MNIST.py

```

1  import numpy as np
2  import sys
3  import os
4  import cntk
5  from cntk.train.training_session import *
6
7  # folder which includes "Train-28x28_cntk_text.txt"
8  data_path = sys.argv[1]
9  # folder for model output
10 model_path = sys.argv[2]
11
12 # Creates and trains a feedforward classification model for MNIS
13 def train_mnist():
14     # variables denoting the features and label data
15     input_var = cntk.input_variable((1, 28, 28))
16     label_var = cntk.input_variable(10)
17
18     # instantiate the feedforward classification model
19     with cntk.layers.default_options(
20         init = cntk.layers.glorot_uniform(),
21         activation=cntk.ops.relu):
22         h = input_var/255.0
23         h = cntk.layers.Convolution2D((5,5), 32, pad=True)(h)
24         h = cntk.layers.MaxPooling((3,3), (2,2))(h)
25         h = cntk.layers.Convolution2D((3,3), 48)(h)
26         h = cntk.layers.MaxPooling((3,3), (2,2))(h)
27         h = cntk.layers.Convolution2D((3,3), 64)(h)
28         h = cntk.layers.Dense(96)(h)
29         h = cntk.layers.Dropout(0.5)(h)
30         z = cntk.layers.Dense(10, activation=None)(h)
31

```

```

32 # create source for training data
33 reader_train = cntk.io.MinibatchSource(
34     cntk.io.CTFDeserializer(
35         os.path.join(data_path, 'Train-28x28_cntk_text.txt'),
36         cntk.io.StreamDefs(
37             features = cntk.io.StreamDef(field='features', shape=28
38             labels = cntk.io.StreamDef(field='labels', shape=10)
39         )),
40     randomize=True,
41     max_samples = 60000,
42     multithreaded_deserializer=True)
43
44 # create trainer
45 epoch_size = 60000
46 minibatch_size = 64
47 max_epochs = 40
48
49 lr_schedule = cntk.learning_rate_schedule(
50     0.2,
51     cntk.UnitType.minibatch)
52 learner = cntk.sgd(
53     z.parameters,
54     lr_schedule)
55 dist_learner = cntk.train.distributed.data_parallel_distribute
56     learner = learner,
57     num_quantization_bits=32, # non-quantized SGD
58     distributed_after=0) # all data is parallelized
59
60 progress_printer = cntk.logging.ProgressPrinter(
61     tag='Training',
62     num_epochs=max_epochs,
63     freq=100,
64     log_to_file=None,
65     rank=cntk.train.distributed.Communicator.rank(),
66     gen_heartbeat=True,
67     distributed_freq=None)
68 ce = cntk.losses.cross_entropy_with_softmax(
69     z,
70     label_var)
71 pe = cntk.metrics.classification_error(
72     z,
73     label_var)
74 trainer = cntk.Trainer(
75     z,
76     (ce, pe),
77     dist_learner,
78     progress_printer)
79
80 # define mapping from reader streams to network inputs
81 input_map = {
82     input_var : reader_train.streams.features,
83     label_var : reader_train.streams.labels
84 }
85
86 cntk.logging.log_number_of_parameters(z) ; print()
87
88 session = training_session(
89     trainer=trainer,
90     mb_source = reader_train,
91     model_inputs_to_streams = input_map,
92     mb_size = minibatch_size,
93     progress_frequency=epoch_size,
94     checkpoint_config = CheckpointConfig(
95         frequency = epoch_size,
96         filename = os.path.join(
97             model_path, "ConvNet_MNIST"),
98         restore = False),
99     test_config = None
100 )
101 session.train()

```

```

102
103 # save model only by one process
104 if 0 == cntk.Communicator.rank():
105     z.save(
106         os.path.join(
107             model_path,
108             "ConvNet_MNIST.dnn"))
109
110 # clean-up distribution
111 cntk.train.distributed.Communicator.finalize();
112
113 return
114
115 if __name__ == '__main__':
116     train_mnist()

```

Train-28x28_cntk_text.txt

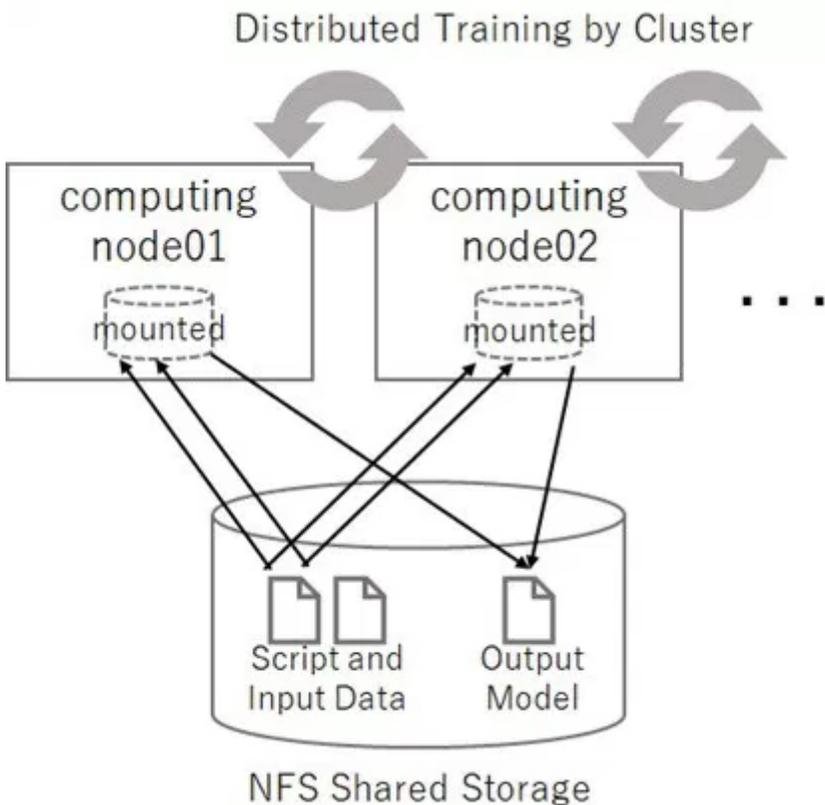
```

|labels 0 0 0 0 0 0 0 1 0 0 |features 0 3 18 18 126... (784 pixel data)
|labels 0 1 0 0 0 0 0 0 0 0 |features 0 0 139 0 253... (784 pixel data)
...

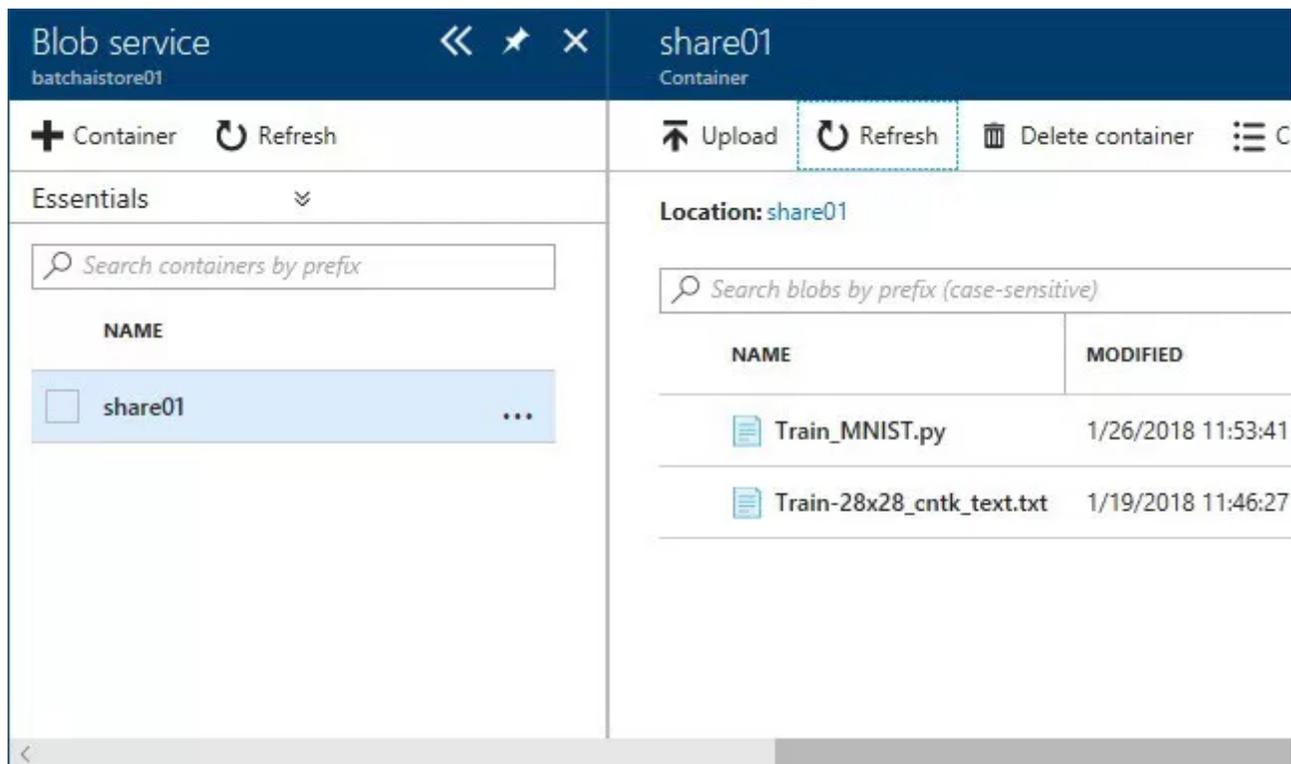
```

Azure Batch AI uses shared script and data (both input data and model output) for the distributed training. Therefore it uses the shared storage like Azure blob, Azure file share, or NFS file servers as following illustrated.

In this post we use Azure Blob Storage and we assume that the storage account name is “batchaistore01” and the container name is “share01”.



Before starting, locate the previous Train_MNIST.py (script) and Train-28x28_cntk_text.txt (input data) in the blob container.



Now let's create the cluster (multiple computing nodes) with the following CLI command.

Here we're creating 2 nodes of Data Science Virtual Machines (DSVM), in which the required software is all pre-configured including GPU-accelerated drivers, OpenMPI and deep learning libraries (TensorFlow, Cognitive Toolkit, Caffe, Keras, etc) with GPU utilized.

Instead of using DSVM, you can also use simple Ubuntu 16.04 LTS (use "UbuntuLTS" instead of "UbuntuDSVM" in the following command) and create a job with the container image, in which the required software is configured. (Later I'll show you how to use the container image in the job.)

Note : When you don't need the distributed training (you only use one single node for the training), you can just specify "--min 1" and "--max 1" in the following command.

Note : Currently (in preview) the supported regions are EastUS, WestUS2, WestEurope and available VM Size (pricing tier) is D-series, E-series, F-series, NC-series, ND-series.

```
# create cluster (computing nodes)
az batchai cluster create --name cluster01 \
  --resource-group testrg01 \
  --location eastus \
  --vm-size STANDARD_NC6 \
  --image UbuntuDSVM \
  --min 2 --max 2 \
  --storage-account-name batchaistore01 \
  --storage-account-key ciUzonyM... \
  --container-name share01 \
  --user-name tsmatsuz --password P@ssw0rd
```

Please remember my [previous post](#), in which we manually configured the computing node.

Using Batch AI, the only one command runs all configuration tasks including MPI and inter-node's communication settings. It significantly simplifies the distributed training !

After you've created the cluster, please wait until "allocationState" is steady without any errors using "az batchai cluster show" command. (See the following results.)

```
# show status for cluster creation
```

```

az batchai cluster show \
  --name cluster01 \
  --resource-group testrg01

{
  "allocationState": "steady",
  "creationTime": "2017-12-18T07:01:00.424000+00:00",
  "currentNodeCount": 1,
  "errors": null,
  "location": "EastUS",

  "name": "cluster01",
  "nodeSetup": {
    "mountVolumes": {
      "azureBlobFileSystems": null,
      "azureFileShares": [
        {
          "accountName": "batchaistore01",
          "azureFileUrl": "https://batchaistore01.file.core.windows.net/share01",
          "credentials": {
            "accountKey": null,
            "accountKeySecretReference": null
          },
          "directoryMode": "0777",
          "fileMode": "0777",
          "relativeMountPath": "dir01"
        }
      ],
      "fileServers": null,
      "unmanagedFileSystems": null
    },
    "setupTask": null
  },
  ...
}

```

You can get the access information for the computing nodes with the following command.

As you can see, here's 2 computing nodes : one is accessed by ssh with port 50000 and another with port 50001. (The inbound NAT is used for accessing nodes with public addresses.)

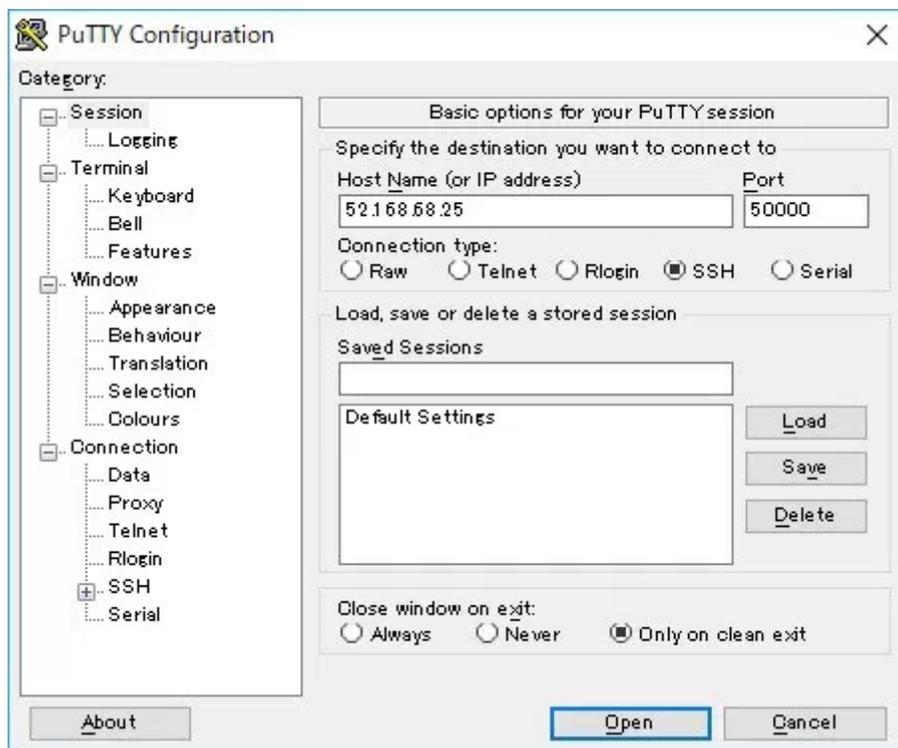
```

# Show endpoints
az batchai cluster list-nodes \
  --name cluster01 \
  --resource-group testrg01

[
  {
    "ipAddress": "52.168.68.25",
    "nodeId": "tvm-4283973576_1-20171218t070031z",
    "port": 50000.0
  },
  {
    "ipAddress": "52.168.68.25",
    "nodeId": "tvm-4283973576_2-20171218t070031z",
    "port": 50001.0
  }
]

```

Please login to the computing node with ssh client, and you can find that the storage container “share01” is mounted as path “\$AZ_BATCHAI_MOUNT_ROOT/bfs” (/mnt/batch/tasks/shared/LS_root/mounts/bfs in my environment) in your computing node. (You can change the mount point “bfs” with the cluster creation option.)



When you set the same value for minimum (min) and maximum (max) as cluster nodes’ count with cluster creation option, the auto-scale setting is disabled and it uses the fixed nodes for the cluster.

When you set the different values, the auto-scale setting is enabled automatically. (See the following result.)

You can check whether the auto-scaling is enabled with “az batchai cluster show” command.

Note : When using the fixed nodes, you can change the node count with “az batchai cluster resize“. When using auto-scaling, you can change the auto-scale configuration with “az batchai cluster auto-scale“.

Note : As you can see, the running workload won’t be scaled out on-the-fly, even if you enabled auto-scaling. Auto-scale settings will effect, when you run multiple workloads on the cluster.

```
# creating fixed nodes...
az batchai cluster create --name cluster01 \
  --resource-group testrg01 \
  --location eastus \
  --min 2 --max 2 ...
```

```
{
  "allocationState": "steady",
  "scaleSettings": {
    "autoScale": null,
    "manual": {
      "nodeDeallocationOption": "requeue",
      "targetNodeCount": 2
    }
  },
  ...
}
```

```
# creating auto-scaling nodes...
az batchai cluster create --name cluster01 \
  --resource-group testrg01 \
  --location eastus \
  --min 1 --max 3 ...

{
  "allocationState": "steady",
  "scaleSettings": {
    "autoScale": {
      "initialNodeCount": 0,
      "maximumNodeCount": 3,
      "minimumNodeCount": 1
    },
    "manual": null
  },
  ...
}
```

When your cluster is ready, now you can publish the training job !

Before starting your job, you must first create the following json for the job definition.

With this definition, the command “python Train_MNIST.py {dir for training data} {dir for model output}” will be launched on MPI (on “mpirun” command). Therefore the training workloads will run on all nodes in parallel and the results are exchanged each other.

Note : When you set 2 as “nodeCount” as follows, 2 computing nodes are set in the host file for MPI (located as \$AZ_BATCHAI_MPI_HOST_FILE) and the file is used by “mpirun” command.

Note : Here we use the settings for CNTK (Cognitive Toolkit), but you can also specify the framework settings for **TensorFlow**, **Caffe** (incl. Caffe2) and **Chainer** with Batch AI.

Later I’ll explain more about job definition file...

job.json

```
{
  "properties": {
    "nodeCount": 2,
    "cntkSettings": {
      "pythonScriptFilePath": "$AZ_BATCHAI_MOUNT_ROOT/bfs/Train_MNIST.py",
      "commandLineArgs": "$AZ_BATCHAI_MOUNT_ROOT/bfs $AZ_BATCHAI_MOUNT_ROOT/bfs/model"
    },
    "stdOutErrPathPrefix": "$AZ_BATCHAI_MOUNT_ROOT/bfs/output"
  }
}
```

You can run this job with the following command.

As I mentioned in my early post (see my post “[Azure Batch – Walkthrough and How it works](#)“), the job runs under an auto-user account named “_azbatch”, and not used the provisioned user account (“tsmatsuz” in my sample).

```
# run job !
az batchai job create --name job01 \
  --cluster-name cluster01 \
  --resource-group testrg01 \
```

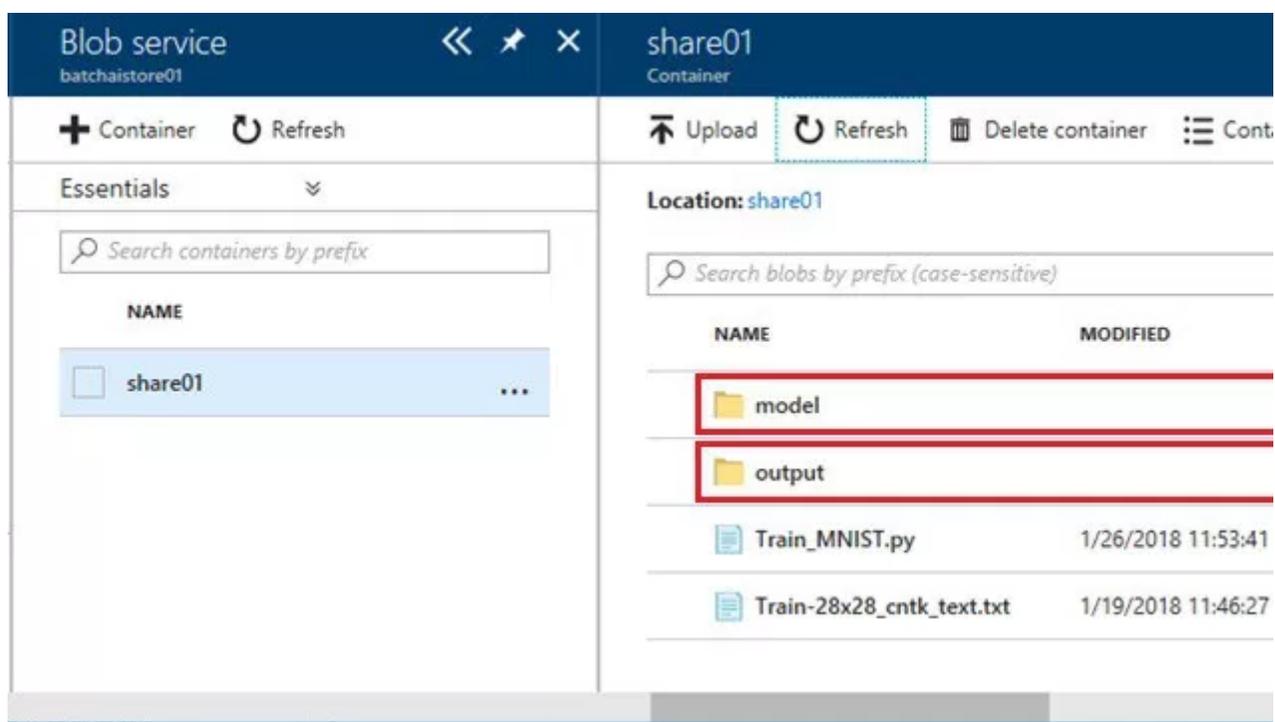
```
--location eastus \  
--config job.json
```

After you published the job, you can see the job processing state (queued, running, succeeded, etc) with the following command.

```
# see the job status  
az batchai job list -o table
```

Name	Resource Group	Cluster	Cluster RG	Tool	Nodes	State	Exit c
job01	testrg01	cluster01	testrg01	cntk	2	succeeded	

The result (the generated model and console output) is located in the shared (mounted) folder and you can download these files. (As I'll explain later, you can also get the download url with CLI command.)



When you want to terminate the queued job, you can run the following command.

```
# terminate job  
az batchai job terminate \  
--name job01 \  
--resource-group testrg01
```

More about job definitions

Here I show you more about job definitions (job.json).

When you use your own favorite framework like Keras, MXNet, etc except for BatchAI-supported frameworks, you can use `customToolkitSettings` in job definitions.

For example, you can define the previous CNTK distributed training using “`customToolkitSettings`” as follows. (The result is the same as before.)

Using custom settings (job.json).

```
{
  "properties": {
    "nodeCount": 2,
    "customToolkitSettings": {
      "commandLine": "mpirun --mca btl_tcp_if_include eth0 --hostfile $AZ_BATCHAI_MPI_H
    },
    "stdoutErrPathPrefix": "$AZ_BATCHAI_MOUNT_ROOT/bfs/output"
  }
}
```

In my previous example, I configured command with only “`$AZ_BATCHAI_MOUNT_ROOT`” in job definition (job.json). But you can also write the definition as follows with good modularity. (Here we’re defining extra variables “`AZ_BATCHAI_INPUT_MYSCRIPT`”, “`AZ_BATCHAI_INPUT_MYDATA`”, and “`AZ_BATCHAI_OUTPUT_MYMODEL`”.)

With inputDirectories and outputDirectories (job.json).

```
{
  "properties": {
    "nodeCount": 2,
    "inputDirectories": [
      {
        "id": "MYSCRIPT",
        "path": "$AZ_BATCHAI_MOUNT_ROOT/bfs"
      },
      {
        "id": "MYDATA",
        "path": "$AZ_BATCHAI_MOUNT_ROOT/bfs"
      }
    ],
    "outputDirectories": [
      {
        "id": "MYMODEL",
        "pathPrefix": "$AZ_BATCHAI_MOUNT_ROOT/bfs",
        "pathSuffix": "model"
      }
    ],
    "cntkSettings": {
      "pythonScriptFilePath": "$AZ_BATCHAI_INPUT_MYSCRIPT/Train_MNIST.py",
      "commandLineArgs": "$AZ_BATCHAI_INPUT_MYDATA $AZ_BATCHAI_OUTPUT_MYMODEL"
    },
    "stdoutErrPathPrefix": "$AZ_BATCHAI_MOUNT_ROOT/bfs/output"
  }
}
```

The benefit of this writing is not only for modularity.

When you write with “`outputDirectories`” as above, you can get the details about output files with this id of “`outputDirectories`” as follows. Here are 3 files (ConvNet_MNIST, ConvNet_MNIST.ckp, ConvNet_MNIST.dnn) in the directory `AZ_BATCHAI_OUTPUT_MYMODEL`.

(When you access job stdout and stderr, use “`stdouterr`” for directory id.)

```
# list files in "MYMODEL"
az batchai job list-files \
  --name job03 \
  --output-directory-id MYMODEL \
  --resource-group testrg01

[
  {
    "contentType": "application/octet-stream",
    "downloadUrl": "https://batchaistore01.blob.core.windows.net/share01/b3a...",
    "name": ".../outputs/model/ConvNet_MNIST"
  },
  {
    "contentType": "application/octet-stream",
    "downloadUrl": "https://batchaistore01.blob.core.windows.net/share01/b3a...",
    "name": ".../outputs/model/ConvNet_MNIST.ckp"
  },
  {
    "contentType": "application/octet-stream",
    "downloadUrl": "https://batchaistore01.blob.core.windows.net/share01/b3a...",
    "name": ".../outputs/model/ConvNet_MNIST.dnn"
  }
]
```

You can also view the results with integrated UI in Azure Portal or AI tools, etc.

Azure Portal

The screenshot shows the Azure Portal interface for a Batch AI job named "job05". The interface includes a navigation pane on the left with sections for Overview, Activity log, Access control (IAM), Tags, SETTINGS (Nodes, Output files, Input directory, Output directory, Locks), and GENERAL (Automation script, Properties). The main content area displays the job's progress through four stages: Create (0.91s), Queue (1.82s), Run (25.31s), and Complete. The execution state is "succeeded" with a total duration of 28.3s, submitted on Monday, January 29, 2018 at 14:52:42. Below this, there is a section for "Output directory" with a dropdown menu showing "MYMODEL" and "stdouterr". Three output files are listed with their respective download URLs.

AI tools for Visual Studio

The screenshot shows the Azure Batch AI Job Browser interface. The main window displays a list of jobs on the left and detailed information for the selected job (job05) on the right. The job is in a 'Succeeded' state. The 'Properties' tab is active, showing the following details:

- StdOutErr Path Prefix:** \$AZ_BATCHAI_MOUNT_ROOT/bfs/output
- Input Directories:**

Id	Path
MYSCRIPT	\$AZ_BATCHAI_MOUNT_ROOT/bfs
MYDATA	\$AZ_BATCHAI_MOUNT_ROOT/bfs
- Output Directories:**

Id	Path Prefix	Path Suffix	Output Type	Create New
MYMODEL	\$AZ_BATCHAI_MOUNT_ROOT/bfs	model	custom	True

In my previous example, we used only built-in modules in python.
If you want to install some additional module, please use the job preparation as follows.

Invoking pre-task (job.json)

```
{
  "properties": {
    "nodeCount": 1,
    "jobPreparation": {
      "commandLine": "pip install mpi4py"
    },
    ...
  }
}
```

You can also specify commands using sh file as follows.

Invoking pre-task (job.json)

```
{
  "properties": {
    "nodeCount": 1,
    "jobPreparation": {
      "commandLine": "bash $AZ_BATCHAI_MOUNT_ROOT/myprepare.sh"
    },
    ...
  }
}
```

As I mentioned earlier, you can use the docker image without DSVM provisioning.
With the following definition, you can use simple Ubuntu VM for cluster creation and use docker image for provisioning.

With docker image (job.json)

```
{
  "properties": {
    "nodeCount": 2,
    "cntkSettings": {
      "pythonScriptFilePath": "$AZ_BATCHAI_MOUNT_ROOT/bfs/Train_MNIST.py",
      "commandLineArgs": "$AZ_BATCHAI_MOUNT_ROOT/bfs $AZ_BATCHAI_MOUNT_ROOT/bfs/model"
    },
    "stdOutErrPathPrefix": "$AZ_BATCHAI_MOUNT_ROOT/bfs/output",
    "containerSettings": {
      "imageSourceRegistry": {
        "image": "microsoft/cntk:2.1-gpu-python3.5-cuda8.0-cudnn6.0"
      }
    }
  }
}
```

Next I show you Azure Distributed Data Engineering Toolkit (aztk), which is also built on top of Azure Batch and provides Spark cluster with docker images.

Enjoy your AI works with Azure !

[Reference]

Github – Azure Batch AI recipes

<https://github.com/Azure/BatchAI/tree/master/recipes>