

IBM Watson ML Community Edition

integration with Spectrum LSF

An abstract graphic consisting of a dense, interconnected network of thin blue lines forming a mesh or wireframe structure. The lines are arranged in a way that creates a sense of depth and movement, resembling a stylized wave or a complex data network. The overall effect is a futuristic, digital aesthetic.

Ing. Florin Manaila
Senior Architect and Inventor
Cognitive Systems (HPC and Deep Learning)
IBM Systems Hardware Europe
Member of the IBM Academy of Technology (AoT)

Frankfurt am Main, June 17, 2019




Observation #1



The Problem




Data Scientist



The Problem



Data Scientist




The Problem



Data Scientist

- Typical Power Consumption: **1,500 W**
- Ambient temperature **28-30°C**
- Very good electrical power circuit is needed to avoid overloading the circuit

Accelerated Workstation



Issues with this operational model



Data Scientists


- Bad Utilization of HW resources
- Deployment into production is not issue-free
- Not so easy to share datasets
- Long time data-preparation
- Long training times
- Datasets are duplicated several times
- GPUs waiting for data to come from the NFS shared storage

Observation #2



Parallel Architecture in Deep Learning


Hardware Architecture



Source: ETH Zurich (T. Ben-Nun and T. Hoefer. Demystifying Parallel and Distributed Deep Learning)

Parallel Architecture in Deep Learning


Training with Single vs. Multiple Nodes



Source: ETH Zurich (T. Ben-Nun and T. Hoefer. Demystifying Parallel and Distributed Deep Learning)

Parallel Architecture in Deep Learning

Communication Layer



Source: ETH Zurich (T. Ben-Nun and T. Hoefler. Demystifying Parallel and Distributed Deep Learning)

Distributed Deep Learning


Goals

The overall goal of distributed deep learning is to

reduce the training time

To this end the primary features:

- Automatic Topology Detection
- Rankfile generation
- Automatic mpirun option handling
- Efficiency in scalability

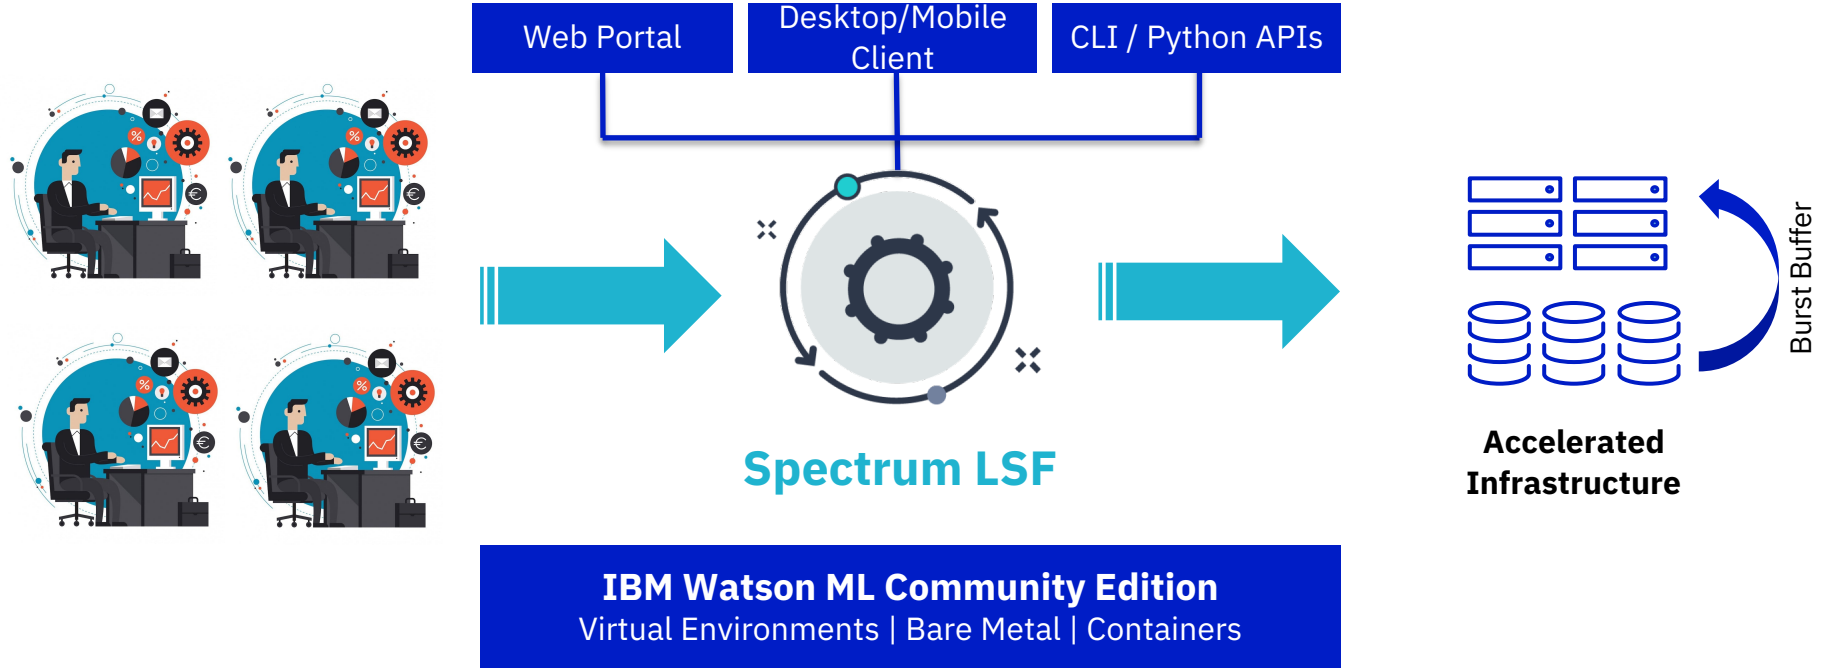


Solution



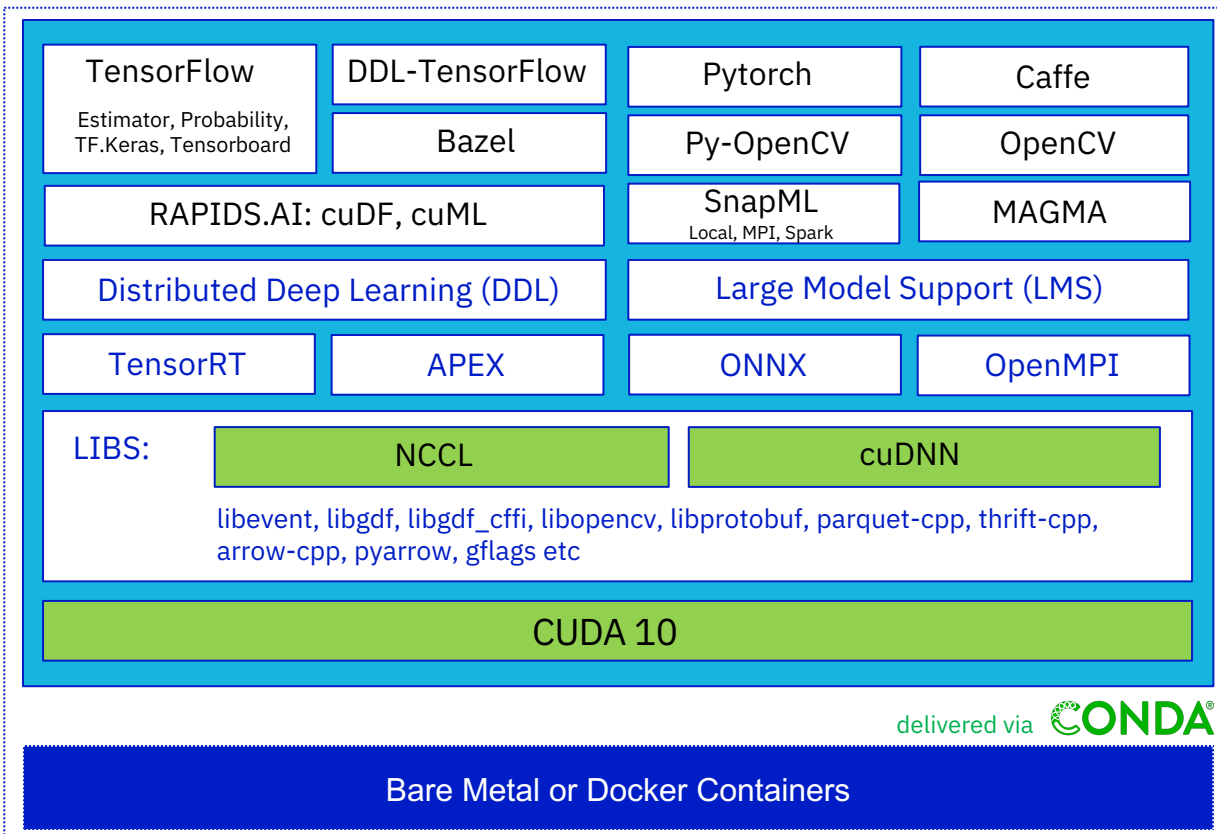
IBM Watson ML Community Edition

integration with with Spectrum LSF





IBM Watson Machine Learning Community Edition





IBM Watson Machine Learning Community Edition

Deep learning training takes
days to weeks

Limited scaling to
multiple x86 servers


PowerAI with DDL enables
scaling to 100s of GPUs

16 Days Down to 7 Hours
58x Faster



ResNet-101, ImageNet-22K

Near Ideal Scaling to 256 GPUs



ResNet-50, ImageNet-1K

Caffe with PowerAI DDL, Running on Minsky (S822Lc) Power System



IBM Watson ML Community Edition

Spectrum LSF Integration

- LSF support has been added to the ddlrn tool
- When using ddlrn from a LSF job, the list of hosts no longer has to be provided. ddlrn will detect the hosts that the job should run on.
- Example:
 - `ddlrn python my_script.py`

NO NEED TO SPECIFY RUNNING HOSTS



IBM Watson ML Community Edition

Spectrum LSF Integration



```
from pythonlsf import lsf

def run_job(command):
    """
    Run a job...
    """
    submitreq = lsf.submit()
    submitreq.command = command
    submitreq.options = 0
    submitreq.options2 = 0

    limits = []
    for i in range(0, lsf.LSF_RLIM_NLIMITS):
        limits.append(lsf.DEFAULT_RLIMIT)

    submitreq.rLimits = limits

    submitreq.beginTime = 0
    submitreq.termTime = 0
    submitreq.numProcessors = 1
    submitreq.maxNumProcessors = 1

    submitreply = lsf.submitReply()

    if lsf.lsb_init("test") > 0:
        exit(1)

    job_id = lsf.lsb_submit(submitreq, submitreply)
    return job_id

if __name__ == '__main__':
    print("LSF Clustername is :", lsf.ls_getclustername())
    print(run_job("/bin/sleep 10"))
```

Spectrum LSF Python Integration

<https://github.com/IBMSpectrumComputing/lsf-python-api>

IBM AC922 Deep Learning Cluster

florin

.....


☐ Keep me logged in




Log in






© Copyright International Business Machines Corp. 1992, 2019. US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp. Portions Copyright © 1991, 1999 Free Software Foundation, Inc.





**IBM AC922**
Deep Learning Cluster

 **Workload**  **System & Settings**  **Reports**

florinux ▾   ▾  1:37:17 AM +0200 | ▾

Workload <


 **New Workload**


 **Workload** ^


By Queue

By Group

Waiting for Input


 **Data**


 **Definitions** v

 **VNC Consoles**

All Submissions


Applications > Deep-Learning


 ☆


Search 

Folders(0)

Applications(3)


MNIST_Training


Python-GENERIC


TensorFlow-Benchmark

^

^

Submission Form: Python-GENERIC

Basic Job Options

Job Name

PowerAI Virtual
Environment Name

powerai_env

Python Script ?

Add Local File

Add Server File

Script Options ?

Advanced

Requirements

Resource
Requirement string

rusage[mem=32GB]

Additional Job Options

Project

ai_training

Submit to this Queue

normal

Other bsub options

-n 10 -gpu 'num=1:mode=exclusive_process' -W 1:

Data

Notification

Submit

Save As

Revert

Close


Name: Python-GENERIC Application: Tensorflow Category: Job Submission

Submission Form Submission Script Help Documentation Data Graphs


Basic Job Options

Job Name

PowerAI Virtual Environment Name

Python Script 

Script Options 

Server 10608B 


Move ... 

Open ... 



Advanced

Requirements

Resource Requirement string 

Additional Job Options

Project

Submit to this Queue

Other bsub options

Save As

Unpublish

Test

Close

Submission Form: MNIST_Training

Basic Job Options

Job Name

How many training steps

10000

Data Directory

./mnist_data

Browse...

Tensorboard ?

☐

PowerAI Virtual Environment Name ?

powerai_env

Advanced

Requirements

Resource Requirement string ?

rusage[mem=5GB]

Additional Job Options

Submit to this Queue

normal

Project

mnist_train

Other bsub options ?

-n 10 -gpu 'num=1:mode=exclusive_process' -W 1:

Data

Submit

Save As

Revert

Close

Submission Form: TensorFlow-Benchmark

Basic Job Options

Job Name	<input type="text"/>	TensorFlow Batch Size	256
TensorFlow No of Batches [?]	200	Number of Warmup Batches [?]	1
Used Data Format	NCHW	TensorFlow Optimizer [?]	sgd
TensorFlow Variable Update [?]	replicated	Number of GPUs for benchmark	1
Single Precision [?]	true	CNN Model	alexnet
PowerAI Virtual Environment Name [?]	powerai_env	Python Version [?]	anaconda3


Advanced

Submit

Save As

Revert

Close



IBM AC922
Deep Learning Cluster

Workload

System & Settings

Reports

florin

2:06:14 AM +0200

Workload

New Workload

Workload

By Queue

By Group

Waiting for Input

Data

Definitions

VNC Consoles

Queue Summary > normal

normal

New

Control

View Output

Delete Directories

Search

User = florin

Ended = Past Hour

<input type="checkbox"/>	ID	Name	State	Application	Submitted	Start Time	Ended	User
<input type="checkbox"/>	292	hpm2	Running	TensorFlow-Ben	2019-04-21 02:0	2019-04-21 02:0	-	florin
<input type="checkbox"/>	291	abc	Done	MNIST_Training	2019-04-21 01:4	2019-04-21 01:4	2019-04-21 01:4	florin
<input type="checkbox"/>	290	hpm1	Done	TensorFlow-Ben	2019-04-21 01:1	2019-04-21 01:1	2019-04-21 01:1	florin

<<

<

Page 1 of 1

>

>>

10


Viewing 1 - 3 of 3




Cognitive Systems Europe / June 17 / © 2019 IBM Corporation








24

Integration Code is Open Sourced

<https://github.com/ticlazau/lsf-integrations>



 **ticlazau / lsf-integrations**
forked from IBMSpectrumComputing/lsf-integrations


 Watch 0  Star 0  Fork 2

 Code  Pull requests 0  Projects 0  Wiki  Security  Insights  Settings









Branch: **master** Create new file Upload files Find file History

[lsf-integrations](#) / [Spectrum LSF Application Center](#) / [TensorflowImage](#) / [submission_templates](#) /

This branch is 20 commits ahead, 33 commits behind IBMSpectrumComputing:master.  Pull request  Compare

 **ticlazau** Delete Python-GENERIC.zip Latest commit 41b886c on Apr 21

..

 Classify_Directory_Of_Images	Adding Tensorflow Image Examples	9 months ago
 Classify_Image	Adding Tensorflow Image Examples	9 months ago
 MNIST_Training	Adding Tensorflow Image Examples	9 months ago
 PowerAI-MNIST_Training	Delete MNIST_Training-GPU.zip	2 months ago
 PowerAI-Python-GENERIC	Delete MNIST_Training_v2.zip	2 months ago
 PowerAI-TensorFlow-Benchmark	Delete Python-GENERIC.zip	2 months ago
 Retrain_Model	Adding Tensorflow Image Examples	9 months ago
 Tensorboard	Adding Tensorflow Image Examples	9 months ago

Python-GENERIC.zip
MNIST_Training.zip
TensorFlow-Benchmark.zip



Notice and disclaimers

Copyright © 2019 by International Business Machines Corporation (IBM). No part of this document may be reproduced or transmitted in any form without written permission from IBM.

U.S. Government Users Restricted Rights — use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.

Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. **This document is distributed “as is” without any warranty, either express or implied. In no event shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity. IBM products and services are warranted according to the terms and conditions of the agreements under which they are provided.**

IBM products are manufactured from new parts or new and used parts. In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply.”

Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.

Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.

Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.

It is the customer's responsibility to insure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer's business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer is in compliance with any law.

Notice and disclaimers cont.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. IBM does not warrant the quality of any third-party products, or the ability of any such third-party products to interoperate with IBM's products. **IBM expressly disclaims all warranties, expressed or implied, including but not limited to, the implied warranties of merchantability and fitness for a particular, purpose.**

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents, copyrights, trademarks or other intellectual property right.

IBM, the IBM logo, ibm.com, AIX, BigInsights, Bluemix, CICS, Easy Tier, FlashCopy, FlashSystem, GDPS, GPFS, Guardium, HyperSwap, IBM Cloud Managed Services, IBM Elastic Storage, IBM FlashCore, IBM FlashSystem, IBM MobileFirst, IBM Power Systems, IBM PureSystems, IBM Spectrum, IBM Spectrum Accelerate, IBM Spectrum Archive, IBM Spectrum Control, IBM Spectrum Protect, IBM Spectrum Scale, IBM Spectrum Storage, IBM Spectrum Virtualize, IBM Watson, IBM Z, IBM z Systems, IBM z13, IMS, InfoSphere, Linear Tape File System, OMEGAMON, OpenPower, Parallel Sysplex, Power, POWER, POWER4, POWER7, POWER8, Power Series, Power Systems, Power Systems Software, PowerHA, PowerLinux, PowerVM, PureApplica- tion, RACF, Real-time Compression, Redbooks, RMF, SPSS, Storwize, Symphony, SystemMirror, System Storage, Tivoli, WebSphere, XIV, z Systems, z/OS, z/VM, z/VSE, zEnterprise and zSecure are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at: www.ibm.com/legal/copytrade.shtml.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both. Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.