



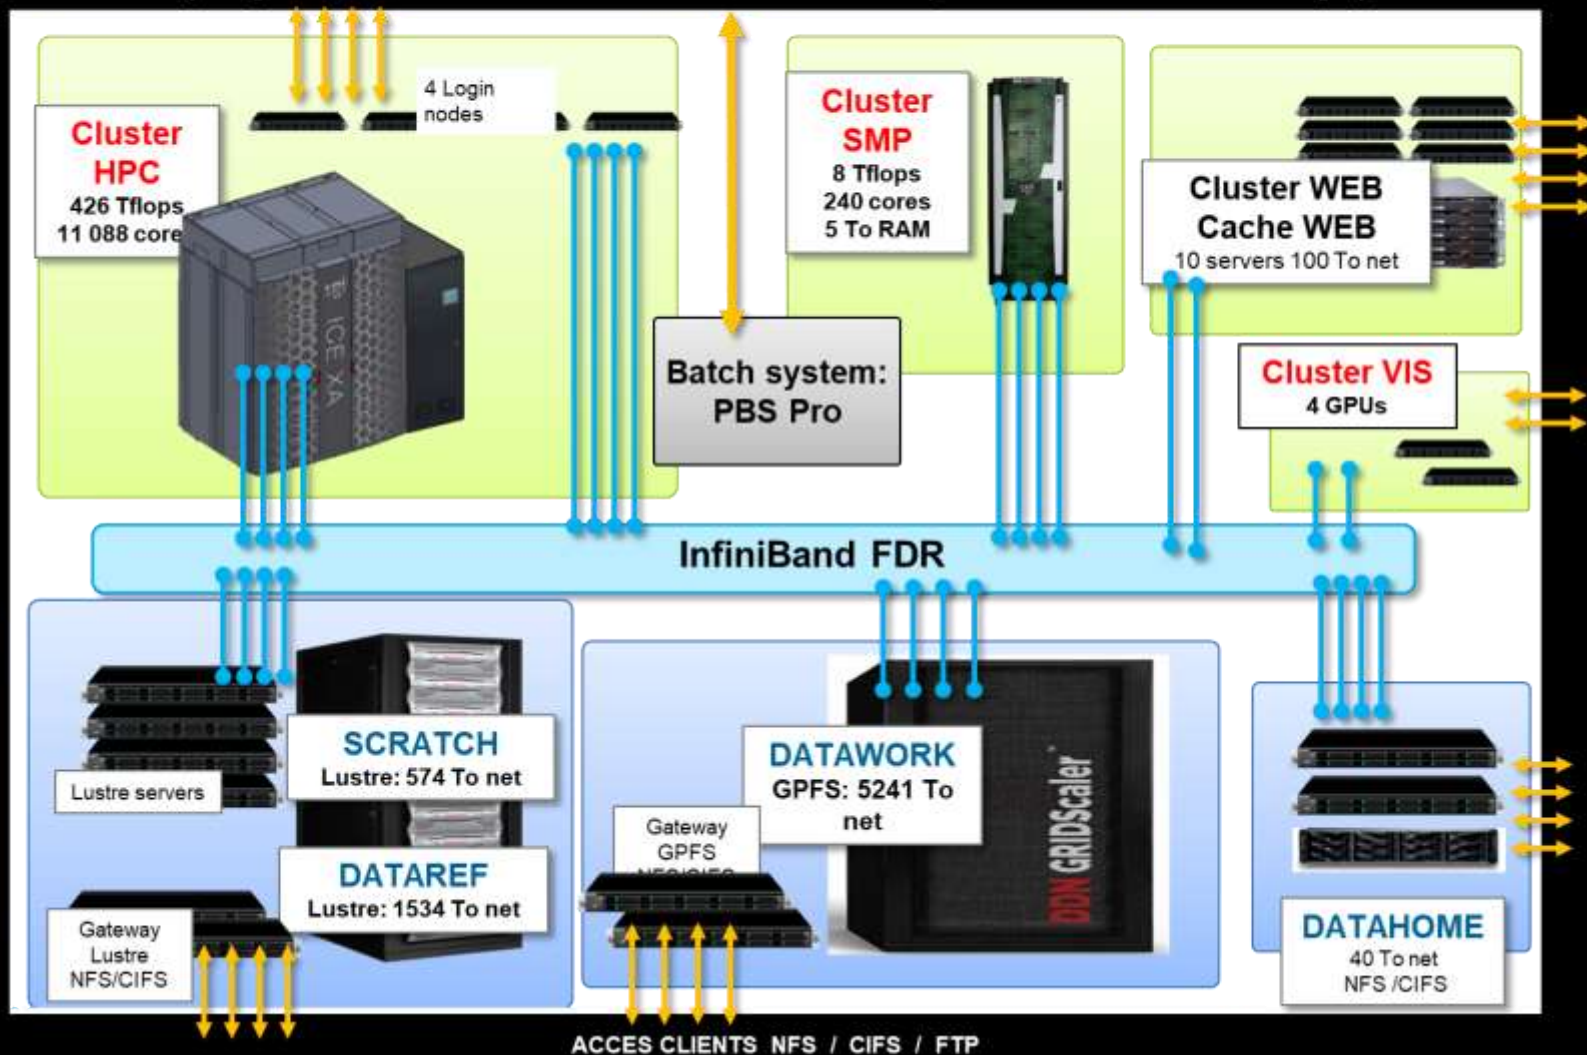
CONFIGURATION OF SPARK/HADOOP ON DATARMOR

Tina Odaka and Antoine QUERIC



DATARMOR

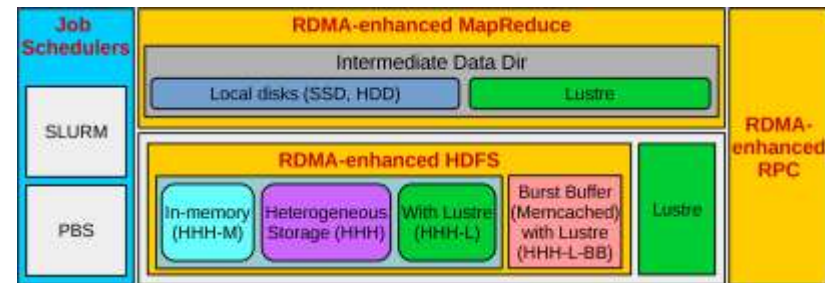
User access
 (10 Gigabit Ethernet connected to IFREMER and INFUSER network (SHOM / IUEM / UBO / ENSTA Bretagne))



DATARMOR is not a 'standard' Hadoop cluster

- Network connections is InfiniBand FDR, 56 times faster than 1Gs 'standard network', and have RDMA (Remote Direct Memory Access) capability.
- Cluster HPC does not have local disk to form 'standard HDFS'.
- DATARMOR have two different parallel file system LUSTRE, and GPFS (IBM Spectrum Scale)

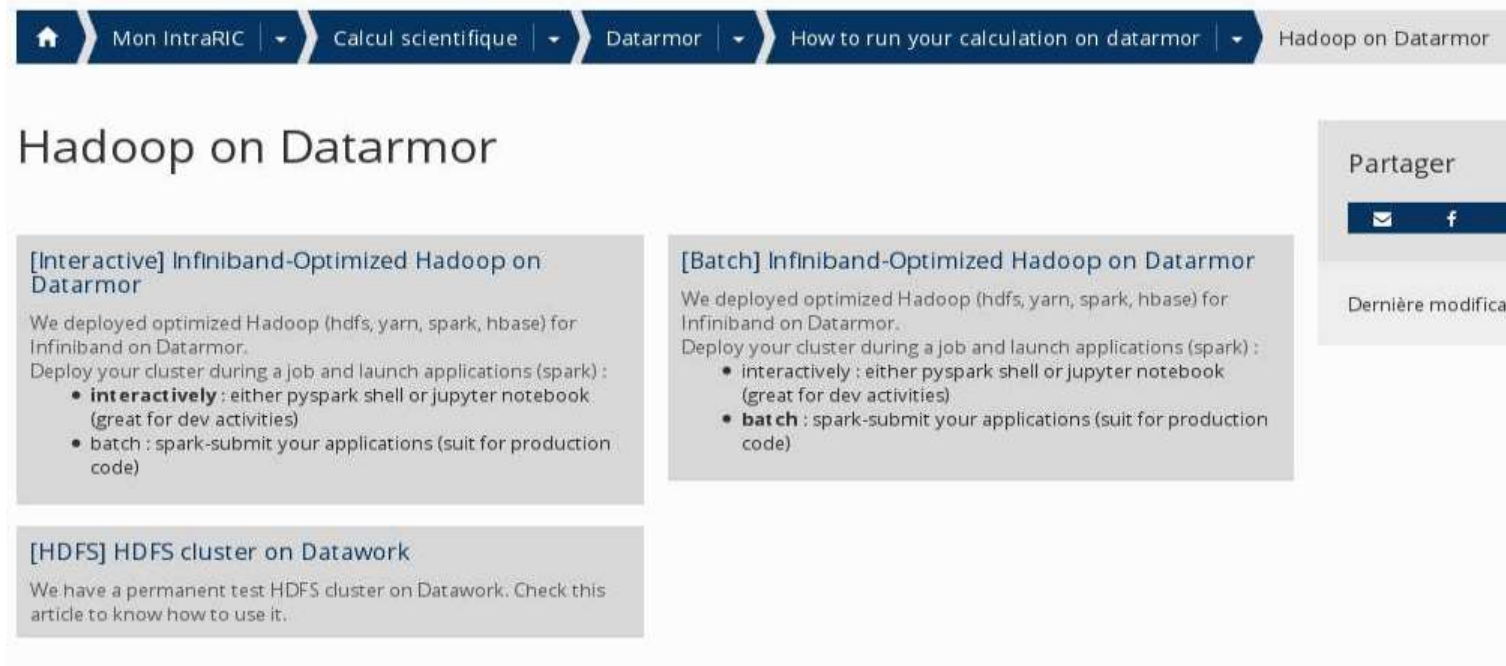
- Infiniband optimized:
RDMA for Apache Spark, Hadoop, and Hbase High-Performance BigData (HiBD)
<http://hibd.cse.ohio-state.edu>



- HDFS transparency
Spectrum scale, IBM product

In practice;

Documentation-> Datarmor –How to run your calculation on datarmor -> Hadoop on Datarmor



Home > Mon IntraRIC > Calcul scientifique > Datarmor > How to run your calculation on datarmor > Hadoop on Datarmor

Hadoop on Datarmor

Partager

✉ f

Dernière modification

[Interactive] Infiniband-Optimized Hadoop on Datarmor

We deployed optimized Hadoop (hdfs, yarn, spark, hbase) for Infiniband on Datarmor.

Deploy your cluster during a job and launch applications (spark) :

- **interactively** : either pyspark shell or jupyter notebook (great for dev activities)
- **batch** : spark-submit your applications (suit for production code)

[Batch] Infiniband-Optimized Hadoop on Datarmor

We deployed optimized Hadoop (hdfs, yarn, spark, hbase) for Infiniband on Datarmor.

Deploy your cluster during a job and launch applications (spark) :

- **interactively** : either pyspark shell or jupyter notebook (great for dev activities)
- **batch** : spark-submit your applications (suit for production code)

[HDFS] HDFS cluster on Datawork

We have a permanent test HDFS cluster on Datawork. Check this article to know how to use it.

In practice

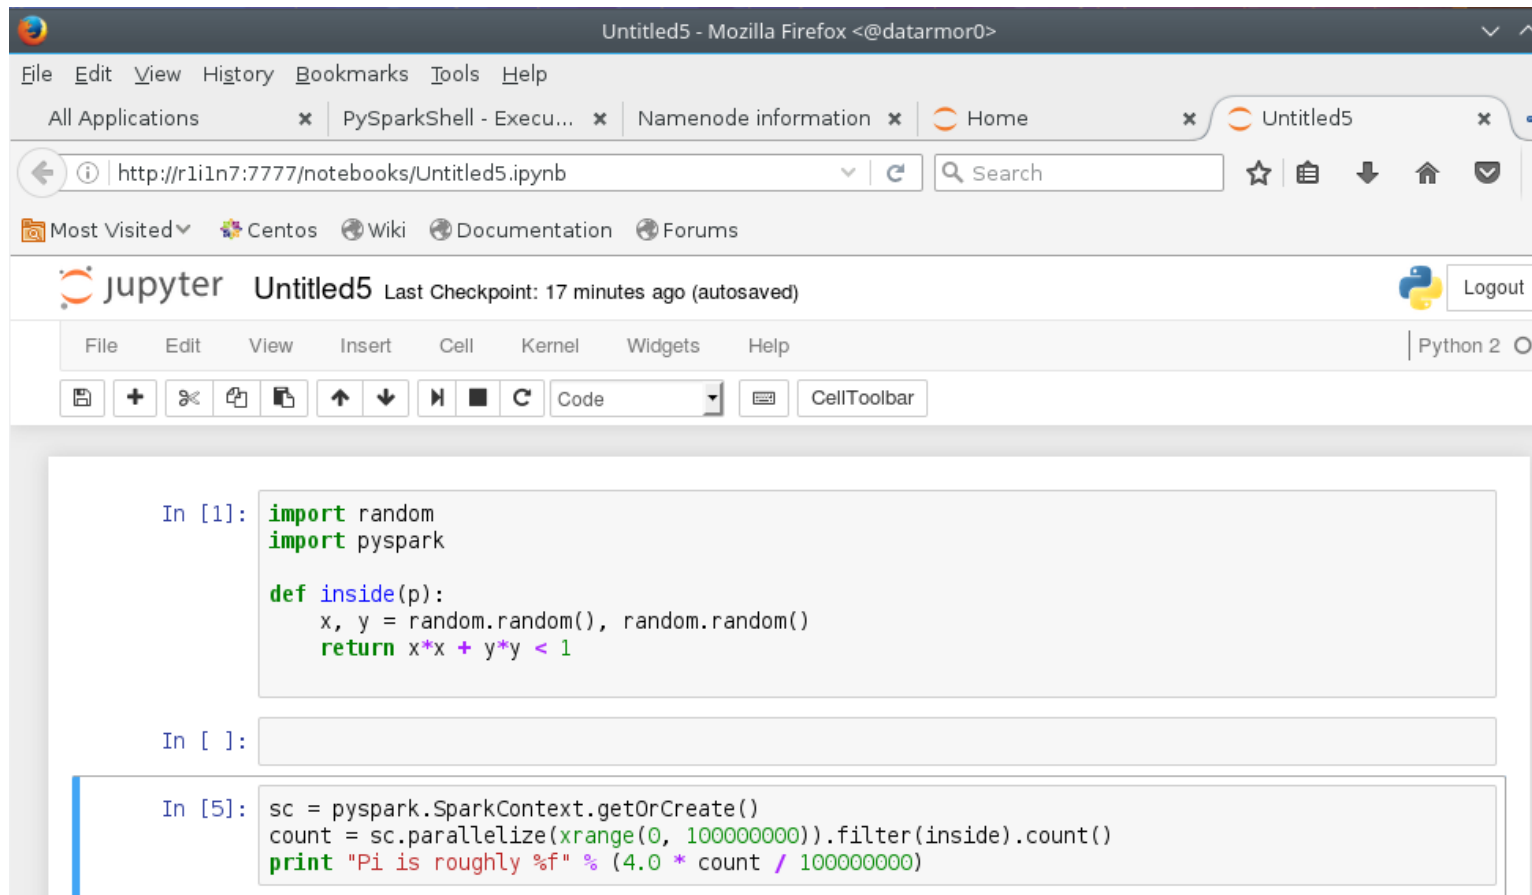
```

Terminal - todaka@vdatarmor-web1:~
File Edit View Terminal Tabs Help
datahome/todaka% qsub -I -q mpi_4 -l walltime=01:00:00
qsub (Warning): Interactive jobs will be treated as not rerunnable
qsub: waiting for job 1115846.datarmor0 to start
qsub: job 1115846.datarmor0 ready

rli0n0 datahome/todaka% bash
todaka@rli0n0:~>
todaka@rli0n0:~> source /usr/share/Modules/3.2.10/init/bash
todaka@rli0n0:~> module load rdma-hadoop/1.3.5
todaka@rli0n0:~> quick-hadoop.sh
Starting HDFS on your Scratch directory !
Setting up Hadoop configuration ...
Your HDFS cluster will write its data on /home1/scratch/todaka/hibd_hadoop/1115846.datarmor0...
Creating files and directories ...
Warning: Permanently added 'rli0n0.ib0.ice.ifremer.fr' (ECDSA) to the list of known hosts.
Formatting HDFS cluster ...
Launch HDFS cluster...
Launch Yarn cluster...
=====
=====
If you want to launch spark applications or issue hdfs commands (eg. ls, cat, rm, put, get ...), please issue first following command to have proper configuration and path :
source /appli/hibd/rdma-hadoop-2.x-1.3.5-x86/sbin/quick-hadoop-get-env.sh --conda-env rdma-hadoop-1.3.5
=====
=====
Your web UI for HDFS / Yarn clusters. Open a new shell window and issue following command :
ssh -X datarmor0-10g firefox --no-remote --new-tab -url http://rli0n0:8088 --new-tab -url http://rli0n0:50070 --new-tab -url http://rli0n0:16010
todaka@rli0n0:~>
todaka@rli0n0:~>
todaka@rli0n0:~>

```

Jupyter notebook: pyspark



Untitled5 - Mozilla Firefox <@datarmor0>

File Edit View History Bookmarks Tools Help

All Applications x PySparkShell - Execu... x Namenode information x Home x Untitled5 x

http://r1i1n7:7777/notebooks/Untitled5.ipynb Search

Most Visited Centos Wiki Documentation Forums

jupyter Untitled5 Last Checkpoint: 17 minutes ago (autosaved) Logout

File Edit View Insert Cell Kernel Widgets Help Python 2

In [1]:

```
import random
import pyspark

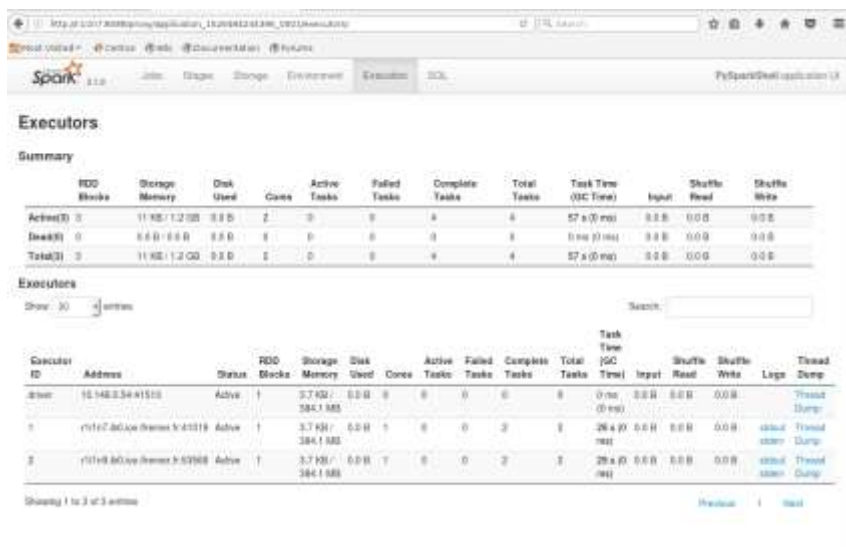
def inside(p):
    x, y = random.random(), random.random()
    return x*x + y*y < 1
```

In []:

In [5]:

```
sc = pyspark.SparkContext.getOrCreate()
count = sc.parallelize(xrange(0, 100000000)).filter(inside).count()
print "Pi is roughly %f" % (4.0 * count / 100000000)
```

spark@datarmor



Executors

Summary

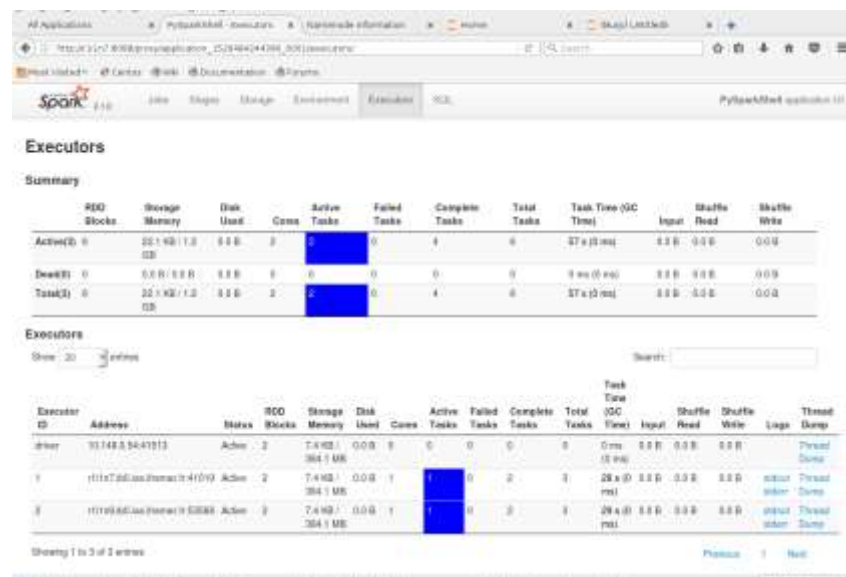
	RDD Blocks	Storage Memory	Disk Used	Cores	Active Tasks	Failed Tasks	Complete Tasks	Total Tasks	Task Time (GC Time)	Input	Shuffle Read	Shuffle Write
Active(3)	3	11 KB / 1.2 GB	0.0 B	2	0	0	4	4	57 s (0 ms)	0.0 B	0.0 B	0.0 B
Deads(0)	0	0.0 B / 0.0 B	0.0 B	0	0	0	0	0	0 ms (0 ms)	0.0 B	0.0 B	0.0 B
Total(3)	3	11 KB / 1.2 GB	0.0 B	2	0	0	4	4	57 s (0 ms)	0.0 B	0.0 B	0.0 B

Executors

Show: 30 entries

Executor ID	Address	Status	RDD Blocks	Storage Memory	Disk Used	Cores	Active Tasks	Failed Tasks	Complete Tasks	Total Tasks	Task Time (GC Time)	Input	Shuffle Read	Shuffle Write	Logs	Thread Dump
driver	10.148.2.34@1513	Active	1	5.7 KB / 384.1 MB	0.0 B	0	0	0	0	0	0 ms (0 ms)	0.0 B	0.0 B	0.0 B		Thread Dump
1	r111728@ip-10-148-2-34	Active	1	3.7 KB / 384.1 MB	0.0 B	1	0	0	2	1	28 s (0 ms)	0.0 B	0.0 B	0.0 B	stdout	Thread Dump
2	r111198@ip-10-148-2-34	Active	1	3.7 KB / 384.1 MB	0.0 B	1	0	0	2	1	28 s (0 ms)	0.0 B	0.0 B	0.0 B	stdout	Thread Dump

Showing 1 to 3 of 3 entries



Executors

Summary

	RDD Blocks	Storage Memory	Disk Used	Cores	Active Tasks	Failed Tasks	Complete Tasks	Total Tasks	Task Time (GC Time)	Input	Shuffle Read	Shuffle Write
Active(0)	0	22.1 KB / 1.2 GB	0.0 B	2	0	0	4	0	57 s (0 ms)	0.0 B	0.0 B	0.0 B
Deads(0)	0	0.0 B / 0.0 B	0.0 B	0	0	0	0	0	0 ms (0 ms)	0.0 B	0.0 B	0.0 B
Total(0)	0	22.1 KB / 1.2 GB	0.0 B	2	0	0	4	0	57 s (0 ms)	0.0 B	0.0 B	0.0 B

Executors

Show: 20 entries

Executor ID	Address	Status	RDD Blocks	Storage Memory	Disk Used	Cores	Active Tasks	Failed Tasks	Complete Tasks	Total Tasks	Task Time (GC Time)	Input	Shuffle Read	Shuffle Write	Logs	Thread Dump
driver	10.148.2.34@1513	Active	2	7.4 KB / 384.1 MB	0.0 B	0	0	0	0	0	0 ms (0 ms)	0.0 B	0.0 B	0.0 B		Thread Dump
1	r111728@ip-10-148-2-34	Active	2	7.4 KB / 384.1 MB	0.0 B	1	0	0	2	1	28 s (0 ms)	0.0 B	0.0 B	0.0 B	stdout	Thread Dump
2	r11198@ip-10-148-2-34	Active	2	7.4 KB / 384.1 MB	0.0 B	1	0	0	2	1	28 s (0 ms)	0.0 B	0.0 B	0.0 B	stdout	Thread Dump

Showing 1 to 3 of 3 entries

status

Now:

Checking the installation with beta, experienced users;

Next step:

benchmark to verify the performance.