

YARN-5983 End-to-End Test Report

Tang, Zhankun

Test Environment	1
FPGA demo application	2
Test scripts	4
Test Cases	5
1. Test FPGA demo application with FPGA and IP env specified	6
2. Test FPGA demo application with FPGA and IP already programmed	8
2.1 Use existing IP Device	8
2.2 Another demo FPGA application with different IP	9
3. Test FPGA demo application with FPGA but no IP env specified	10
4. Test demo application with FPGA but no IP file uploaded	10
5. Test demo application without FPGA devices requested	11
6. No fpga resource plugin configuration	13
Untested End-to-End Scenarios	14

Test Environment

Hardware environment	A desktop with: Intel(R) Core(TM) i7-6700 1 Nallatech 385a card
Software environment	RedHat 6.8 Intel FPGA SDK for OpenCL Nallatech BSP Hadoop 3.1.0-SNAPSHOT single node cluster Modified Distributed-Shell requesting FPGA Modified Spark requesting FPGA

We mainly use the modified Distributed-Shell for below test cases. We added a “-aocx” option and “-container_fpga_count” to make the DS client can upload the IP file and change its container request with a FPGA devices count.

```
[yarn@fpga-devkit apache-hadoop-install-dir]$ yarn jar home/share/hadoop/yarn/hadoop-yarn-applications-distributedshell-3.1.0-SNAPSHOT.jar
WARNING: YARN LOG DIR has been replaced by HADOOP_LOG DIR. Using value of YARN_LOG DIR.
2017-11-25 00:31:22,214 INFO distributedshell.Client: ZHANKUN: add aocx and container_fpga_count options
2017-11-25 00:31:22,215 INFO distributedshell.Client: Initializing Client, Zhankun
No args specified for client to initialize
usage: Client
-aocx <arg>                aocx file path for FPGA
                             testing----
-appname <arg>              Application Name. Default
                             value - DistributedShell
-attempt_failures_validity_interval <arg> when
                             attempt failures validity
                             interval in milliseconds
                             is set to > 0, the failure
                             number will not take
                             failures which happen out
                             of the validityInterval
                             into failure count. If
                             failure count reaches to
                             maxAppAttempts, the
                             application will be
                             failed.
-container_fpga_count <arg> FPGA devices in container
                             request
-container_max_retries <arg> If container could retry,
                             it specifies max retries
-container_memory <arg>      Amount of memory in MB to
                             be requested to run the
                             shell command
-container_resource_profile <arg> Resource profile for the
                             shell command
-container_retry_error_codes <arg> When retry policy is set
                             to
```

Key code changes:

```
if (0 != fpgacount) {
    resourceCapability.setResourceValue(ResourceInformation.FPGA_URI,
fpgacount);
}
```

FPGA demo application

First of all, I'll introduce the FPGA demo application used for testing. The demo is an example in Nallatech BSP(board support package) doing matrix multiplication with FPGA devices. The key steps of it's main logic related to Hadoop YARN design is that it will find the .aocx file in local directory and load it.

```
// Create the program for all device. Use the first device as the
// representative device (assuming all device are of the same type).
std::string binary_file = getBoardBinaryFile("matrix_mult", device[0]);
printf("Using AOXC: %s\n", binary_file.c_str());
program = createProgramFromBinary(context, binary_file.c_str(), device, num_devices);

// Build the program that was just created.
status = clBuildProgram(program, 0, NULL, "", NULL, NULL);
checkError(status, "Failed to build program");
```

Actually, all OpenCL written application should load the aocx files using above OpenCL APIs. So that's why the IntelFpgaOpenclPlugin in YARN should require user upload .aocx file and find it from container's directory. If user don't do so, we assume the FPGA application can find the aocx file itself.

In this end-to-end testing, we don't change DS to upload this demo's "host" binary and run it from container directory because this location of "host" doesn't affect the test result. For simplicity, we invoke the demo with a script "host.sh" in container to avoid penty of environment setting of IntelFPGASDKForOpenCL:

```
[yarn@fpga-devkit apache-hadoop-install-dir]$ cat
/home/fpga/nallatech-385a/examples_p385a_sch_ax115/matrix_mult/matrix_mult/bin/host.sh
#!/bin/bash
echo "FPGA test:\n"
source /home/yarn/.bashrc
/home/fpga/nallatech-385a/examples_p385a_sch_ax115/matrix_mult/matrix_mult/bin/host
```

```
[yarn@fpga-devkit apache-hadoop-install-dir]$ cat ~/.bashrc
# .bashrc
```

```
# Source global definitions
if [ -f /etc/bashrc ]; then
    . /etc/bashrc
fi
```

```
# User specific aliases and functions
export JAVA_HOME=/usr/java/jdk1.8.0_131
source /home/yarn/apache-hadoop-install-dir/env.sh
#FPGA related
export ALTERAOCLSDKROOT="/home/fpga/intelFPGA_pro/17.0/hld"
export
AOCL_BOARD_PACKAGE_ROOT="/home/fpga/intelFPGA_pro/17.0/hld/board/nalla_pcie"
export QSYS_ROOTDIR="/home/fpga/intelFPGA_pro/17.0/qsys/bin"
export
PATH=$ALTERAOCLSDKROOT/bin:/home/fpga/intelFPGA_pro/17.0/quartus/bin:/home/fpga/intelFPGA_pro/17.0/quartus/sopc_builder/bin:$PATH
export
LD_LIBRARY_PATH=$ALTERAOCLSDKROOT/host/linux64/lib:$AOCL_BOARD_PACKAGE_ROOT/linux64/lib:$LD_LIBRARY_PATH
```

The normal output of demo is as follows:

```
[yarn@fpga-devkit apache-hadoop-install-dir]$ /home/fpga/nallatech-385a/examples_p385a_sch_ax115/matrix_mult/matrix_mult/bin/host.sh
FPGA test:\n
Matrix sizes:
  A: 2048 x 1024
  B: 1024 x 1024
  C: 2048 x 1024
Initializing OpenCL
Platform: Intel(R) FPGA SDK for OpenCL(TM)
Using 1 device(s)
  p385a_sch_ax115 : nalla_pcie (aclnalla_pcie0)
Using AOXC: matrix_mult.aocx
Reprogramming device [0] with handle 1
Generating input matrices
Launching for device 0 (global size: 1024, 2048)

Time: 37.064 ms
Kernel time (device 0): 36.991 ms

Throughput: 115.88 GFLOPS

Computing reference output
Verifying
Verification: PASS
```

Test scripts

There's five scripts created for running DS application for different behaviors.

Fistly, a normal DS with shell command "env" but no FPGA requested:

```
[yarn@fpga-devkit apache-hadoop-install-dir]$ cat test_distributedshell.sh
```

```
#!/bin/bash
```

```
yarn jar
```

```
/home/yarn/apache-hadoop-install-dir/home/share/hadoop/yarn/hadoop-yarn-applications-distributedshell-3.1.0-SNAPSHOT.jar -jar
```

```
/home/yarn/apache-hadoop-install-dir/home/share/hadoop/yarn/hadoop-yarn-applications-distributedshell-3.1.0-SNAPSHOT.jar -shell_command "env"
```

Secondly, run FPGA demo with FPGA devices requested also with IP_ID set in environment:

```
[yarn@fpga-devkit apache-hadoop-install-dir]$ cat
```

```
test_distributedshell_with_FPGA_with_IP.sh
```

```
#!/bin/bash
```

```
yarn jar
```

```
/home/yarn/apache-hadoop-install-dir/home/share/hadoop/yarn/hadoop-yarn-applications-distributedshell-3.1.0-SNAPSHOT.jar -jar
```

```
/home/yarn/apache-hadoop-install-dir/home/share/hadoop/yarn/hadoop-yarn-applications-distributedshell-3.1.0-SNAPSHOT.jar -shell_command
```

```
"/home/fpga/nallatech-385a/examples_p385a_sch_ax115/matrix_mult/matrix_mult/bin/host.sh" -aocx
```

```
/home/fpga/nallatech-385a/examples_p385a_sch_ax115/matrix_mult/matrix_mult/bin/matrix_mult.aocx -shell_env REQUESTED_FPGA_IP_ID=matrix_mult
```

```
-container_fpga_count 1
```

Thirdly, run FPGA demo with FPGA devices requested but no IP_ID set in environment:

```
[yarn@fpga-devkit apache-hadoop-install-dir]$ cat
```

```
test_distributedshell_with_FPGA_without_IP.sh
```

```
#!/bin/bash
yarn jar
/home/yarn/apache-hadoop-install-dir/home/share/hadoop/yarn/hadoop-yarn-applica
tions-distributedshell-3.1.0-SNAPSHOT.jar -jar
/home/yarn/apache-hadoop-install-dir/home/share/hadoop/yarn/hadoop-yarn-applica
tions-distributedshell-3.1.0-SNAPSHOT.jar -shell_command
"/home/fpga/nallatech-385a/examples_p385a_sch_ax115/matrix_mult/matrix_mult/bi
n/host.sh" -aocx
/home/fpga/nallatech-385a/examples_p385a_sch_ax115/matrix_mult/matrix_mult/bi
n/matrix_mult.aocx -container_fpga_count 1
```

Fourthly, run FPGA demo without FPGA devices without IP_ID environment:

```
[yarn@fpga-devkit apache-hadoop-install-dir]$ cat
test_distributedshell_without_FPGA_without_IP.sh
#!/bin/bash
yarn jar
/home/yarn/apache-hadoop-install-dir/home/share/hadoop/yarn/hadoop-yarn-applica
tions-distributedshell-3.1.0-SNAPSHOT.jar -jar
/home/yarn/apache-hadoop-install-dir/home/share/hadoop/yarn/hadoop-yarn-applica
tions-distributedshell-3.1.0-SNAPSHOT.jar -shell_command
"/home/fpga/nallatech-385a/examples_p385a_sch_ax115/matrix_mult/matrix_mult/bi
n/host.sh"
```

Fifthly, run FPGA demo without FPGA devices but set the IP_ID environment:

```
[yarn@fpga-devkit apache-hadoop-install-dir]$ cat
test_distributedshell_without_FPGA_with_IP.sh
#!/bin/bash
yarn jar
/home/yarn/apache-hadoop-install-dir/home/share/hadoop/yarn/hadoop-yarn-applica
tions-distributedshell-3.1.0-SNAPSHOT.jar -jar
/home/yarn/apache-hadoop-install-dir/home/share/hadoop/yarn/hadoop-yarn-applica
tions-distributedshell-3.1.0-SNAPSHOT.jar -shell_command
"/home/fpga/nallatech-385a/examples_p385a_sch_ax115/matrix_mult/matrix_mult/bi
n/host.sh" -aocx
/home/fpga/nallatech-385a/examples_p385a_sch_ax115/matrix_mult/matrix_mult/bi
n/matrix_mult.aocx -shell_env REQUESTED_FPGA_IP_ID=matrix_mult
```

Test Cases

1. Test FPGA demo application with FPGA and IP env specified

YARN configuration for FPGA:

```
<property>
  <description>
    Enable additional discovery/isolation of resources on the NodeManager,
    split by comma. By default, this is empty.
    Acceptable values: { "yarn-io/gpu", "yarn-io/fpga" }.
  </description>
  <name>yarn.nodemanager.resource-plugins</name>
  <value>yarn.io/fpga</value>
</property>

<property>
  <description>
    When yarn.nodemanager.resource.fpga.allowed-fpga-devices=auto specified,
    YARN NodeManager needs to run FPGA discovery binary (now only support
    Intel's aocl) to get FPGA information.
    When value is empty (default), YARN NodeManager will try to locate
    discovery executable from system environment "ALTERAOCLSDKROOT"
  </description>
  <name>yarn.nodemanager.resource-plugins.fpga.path-to-discovery-executables</name>
  <value>/home/fpga/intelFPGA_pro/17.0/hld/bin/aocl</value>
</property>

<property>
  <description>
    Specify FPGA devices which can be managed by YARN NodeManager, split by comma
    Number of FPGA devices will be reported to RM to make scheduling decisions.
    Set to auto (default) let YARN automatically discover FPGA resource from
    system.

    Manually specify FPGA devices if admin only want subset of FPGA devices managed by YARN.
    At present, since we can only configure one major number in c-e.cfg, FPGA device is
    identified by their minor device number. A common approach to get minor
    device number of FPGA is using "aocl diagnose" and check uevent with device name.
  </description>
  <name>yarn.nodemanager.resource-plugins.fpga.allowed-fpga-devices</name>
  <value>auto</value>
</property>

<property>
  <name>yarn.nodemanager.resource-plugins.fpga.vendor-plugin.class</name>
  <value>org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpenclPlugin</value>
</property>
```

```
[yarn@fpga-devkit apache-hadoop-install-dir]$ cat conf/resource-types.xml
<configuration>
  <property>
    <name>yarn.resource-types</name>
    <value>yarn.io/fpga</value>
  </property>
</configuration>
```






```
[yarn@fpga-devkit apache-hadoop-install-dir]$ cat conf/node-resources.xml
<configuration>
  <property>
    <name>yarn.nodemanager.resource.yarn-io/fpga</name>
    <value>1</value>
  </property>
</configuration>
[yarn@fpga-devkit apache-hadoop-install-dir]$
```

FPGA resource initialization log:


```

2017-11-25 09:23:03.163 INFO org.apache.hadoop.yarn.server.nodemanager.NodeManager: registered UNIX signal handlers for [TERM, HUP, INT]
2017-11-25 09:23:03.622 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.FpgaResourcePlugin: Using FPGA vendor plugin: org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpenCLPlugin
2017-11-25 09:23:03.670 WARN org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpenCLPlugin: Failed to execute /home/fpga/intelFPGA_pro/17.0/hld/bin/aocl diagnose, exception message: output:null, continue ...
2017-11-25 09:23:03.670 DEBUG org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpenCLPlugin: aocl diagnose: Running diagnose from /home/fpga/intelFPGA_pro/17.0/hld/board/nalla_pcie/linux64/libexec
----- acl0 -----
Vendor: Nallatech ltd
Phys Dev Name Status Information
aclnalla_pcie0Passed nalla_pcie (aclnalla_pcie0)
PCIe dev_id = 2494, bus:slot.func = 02:00.00, Gen3 x8
FPGA temperature = 52.4 degrees C.
Total Card Power Usage = 31.5 Watts.
Device Power Usage = 0.0 Watts.
DIAGNOSTIC_PASSED
-----
2017-11-25 09:23:03.670 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.FpgaDiscoverer: Trying to diagnose FPGA information ...
2017-11-25 09:23:03.712 WARN org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpenCLPlugin: Failed to execute /home/fpga/intelFPGA_pro/17.0/hld/bin/aocl diagnose, exception message: output:null, continue ...
2017-11-25 09:23:03.716 DEBUG org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpenCLPlugin: aocl diagnose: Running diagnose from /home/fpga/intelFPGA_pro/17.0/hld/board/nalla_pcie/linux64/libexec
----- acl0 -----
Vendor: Nallatech ltd
Phys Dev Name Status Information
aclnalla_pcie0Passed nalla_pcie (aclnalla_pcie0)
PCIe dev_id = 2494, bus:slot.func = 02:00.00, Gen3 x8
FPGA temperature = 52.4 degrees C.
Total Card Power Usage = 31.9 Watts.
Device Power Usage = 0.0 Watts.
DIAGNOSTIC_PASSED
-----
2017-11-25 09:23:03.741 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.CGroupsHandlerImpl: CGroup controller already mounted at: /cgroup/cpu
2017-11-25 09:23:03.741 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.CGroupsHandlerImpl: Initializing mounted controller cpu at /cgroup/cpu/yarn
2017-11-25 09:23:03.750 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.fpga.FpgaResourceHandlerImpl: FPGA Plugin bootstrap success.
2017-11-25 09:23:04.228 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.FpgaNodeResourceUpdateHandler: Initializing configured FPGA resources for the NodeManager.
2017-11-25 09:23:04.228 INFO org.apache.hadoop.yarn.server.nodemanager.NodeStatusUpdaterImpl: NodeManager resources is set to: <memory:8192, vCores:8, yarn.io/fpga: 1>
2017-11-25 09:23:03.794 DEBUG org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpenCLPlugin: Get FPGA major-minor numbers from /dev/aclnalla_pcie0
2017-11-25 09:23:03.795 DEBUG org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpenCLPlugin: stat output:f7:0
2017-11-25 09:23:03.796 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.fpga.FpgaResourceAllocator: Add a list of FPGA Devices: [FPGA Device:(Type: IntelOpenCL, Major: 247, Minor: 0, IPID: null)]

```

Test command	Expectation	Actual Result
start-yarn.sh	No error of NM	
test_distributedshell_with_FPGA_with_IP.sh (-container_fpga_count 1)	A. FPGA allocation log shown	
	B. FPGA IP download log shown	
	C. FPGA IP programming log shown	
	D. Demo pass log shown	

The A, B, C, D log is as follows:

A. For FPGA allocation, the AM(deny FPGA device) should be different with container

```

2017-11-25 09:34:42.992 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.fpga.FpgaResourceHandlerImpl: container_1511626981466_0002_01_000001 requested 0 Intel FPGA(s)
2017-11-25 09:34:42.992 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.fpga.FpgaResourceHandlerImpl: FpgaAllocation:
FpgaAllocation:
Allowed:
Denied
FPGA Device:(Type: IntelOpenCL, Major: 247, Minor: 0, IPID: null)
2017-11-25 09:34:48.306 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.fpga.FpgaResourceHandlerImpl: container_1511626981466_0002_01_000002 requested 1 Intel FPGA(s)
2017-11-25 09:34:48.306 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.fpga.FpgaResourceHandlerImpl: FpgaAllocation:
FpgaAllocation:
Allowed:
Denied
FPGA Device:(Type: IntelOpenCL, Major: 247, Minor: 0, IPID: null)

```

B. For FPGA IP download

```

2017-11-25 09:34:48,308 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpenCLPlugin: Got environment: matrix_mult, search IP file in local
ized resources
2017-11-25 09:34:48,308 DEBUG org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpenCLPlugin: Check: /tmp/hadoop-yarn/nm-local-dir/usercache/yarn/a
ppache/application_1511626981466_0002/filecache/12/matrix_mult.aocx
2017-11-25 09:34:48,308 DEBUG org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpenCLPlugin: Found: /tmp/hadoop-yarn/nm-local-dir/usercache/yarn/
appache/application_1511626981466_0002/filecache/12/matrix_mult.aocx
2017-11-25 09:34:48,308 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.fpga.FpgaResourceHandlerImpl: IP file path: /tmp/hadoop-yarn/nm-local-dir/usercac
he/yarn/appache/application_1511626981466_0002/filecache/12/matrix_mult.aocx

```

C. For FPGA IP programming

```

2017-11-25 09:34:50,561 DEBUG org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpenCLPlugin: aocl program: Running program from /home/fpga/intelF
PGA_pro/17.0/htd/board/nalla_pcie/linux64/libexec
Programming device: p385a_sch_axl15 : nalla_pcie (aclnalla_pcie0)
Program succeed.
2017-11-25 09:34:50,561 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpenCLPlugin: Intel aocl program /tmp/hadoop-yarn/nm-local-dir/user
cache/yarn/appache/application_1511626981466_0002/filecache/12/matrix_mult.aocx to aocl successfully
2017-11-25 09:34:50,561 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.fpga.FpgaResourceAllocator: Update IPID to matrix_mult for this allocated device
:FPGA Device:(Type: IntelOpenCL, Major: 247, Minor: 0, IPID: matrix_mult)

```

D. demo execution pass

```

Container: container_1511626981466_0002_01_000002 on fpga-devkit.sh.intel.com_45360
LogAggregationType: AGGREGATED
=====
LogType:stdout
LogLastModifiedTime:Sat Nov 25 09:34:58 -0700 2017
LogLength:446
LogContents:
FPGA test:\n
Matrix sizes:
  A: 2048 x 1024
  B: 1024 x 1024
  C: 2048 x 1024
Initializing OpenCL
Platform: Intel(R) FPGA SDK for OpenCL(TM)
Using 1 device(s)
  p385a_sch_axl15 : nalla_pcie (aclnalla_pcie0)
Using AOCX: matrix_mult.aocx
Generating input matrices
Launching for device 0 (global size: 1024, 2048)

Time: 37.043 ms
Kernel time (device 0): 36.995 ms

Throughput: 115.95 GFLOPS

Computing reference output
Verifying
Verification: PASS

End of LogType:stdout
*****





```

2. Test FPGA demo application with FPGA and IP already programmed

The YARN configuration is same with case 1.

2.1 Use existing IP Device

Test command	Expectation	Actual Result
--------------	-------------	---------------

test_distributedshell_with_FPGA_without_IP.sh (-container_fpga_count 1) Run twice. Check the second one.	FPGA allocation log	
	FPGA IP download log	
	No FPGA IP programming log	
	Demo pass	

```

2017-11-25 23:03:15,098 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.fpga.FpgaResourceHandlerImpl: container_1511675189604_0009_01_000002 requested 1 Intel FPGA(s)
2017-11-25 23:03:15,098 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.fpga.FpgaResourceHandlerImpl: FpgaAllocation
FpgaAllocation
Allowed:
FPGA Device:(Type: IntelOpenCL, Major: 247, Minor: 0, IPID: matrix_mult)
Denied
2017-11-25 23:03:15,099 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpenCLPlugin: Got environment: matrix_mult, search IP file in local
ized resources
2017-11-25 23:03:15,099 DEBUG org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpenCLPlugin: Check: /tmp/hadoop-yarn/nm-local-dir/usercache/yarn/a
ppcache/application_1511675189604_0009/filecache/12/matrix_mult.aocx
2017-11-25 23:03:15,099 DEBUG org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpenCLPlugin: Found: /tmp/hadoop-yarn/nm-local-dir/usercac
he/yarn/appcache/application_1511675189604_0009/filecache/12/matrix_mult.aocx
2017-11-25 23:03:15,099 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.fpga.FpgaResourceHandlerImpl: IP file path: /tmp/hadoop-yarn/nm-local-dir/usercac
he/yarn/appcache/application_1511675189604_0009/filecache/12/matrix_mult.aocx
2017-11-25 23:03:15,099 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.fpga.FpgaResourceHandlerImpl: IP already in device "acl0,247:0", skip reprogramm
ing

```

2.2 Another demo FPGA application with different IP

Here we run another hello_world FPGA application to show the IPID metadata changing:

```
[yarn@fpga-devkit apache-hadoop-install-dir]$ cat
```

```
test_distributedshell_with_FPGA_with_IP2.sh
```

```
#!/bin/bash
```

```
yarn jar
```

```
/home/yarn/apache-hadoop-install-dir/home/share/hadoop/yarn/hadoop-yarn-applica
tions-distributedshell-3.1.0-SNAPSHOT.jar -jar
```




```
/home/yarn/apache-hadoop-install-dir/home/share/hadoop/yarn/hadoop-yarn-applica
tions-distributedshell-3.1.0-SNAPSHOT.jar -shell_command
```


```
"/home/fpga/nallatech-385a/examples_p385a_sch_ax115/hello_world/hello_world/bi
n/host.sh" -aocx
```

```
/home/fpga/nallatech-385a/examples_p385a_sch_ax115/hello_world/hello_world/bin
```

```
/hello_world.aocx -shell_env REQUESTED_FPGA_IP_ID=hello_world
```

```
-container_fpga_count 1
```

Test command	Expectation	Actual Result
test_distributedshell_with_FPGA_with_IP2.sh (-container_fpga_count 1)	FPGA allocation log	
	FPGA IP download log	
	FPGA IP programming log	

Run test_distributedshell_w ith_FPGA_without_IP.s h first. Check the second.	Demo pass	
--	-----------	---

```

2017-11-25 23:17:52,882 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.fpga.FpgaResourceHandlerImpl: FpgaAllocation:
FpgaAllocation
  Allowed:
    FPGA Device:(Type: IntelOpenCL, Major: 247, Minor: 0, IPID: matrix_mult)
  Denied





2017-11-25 23:17:52,883 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpencilPlugin: Got environment: hello_world, search IP file in local
ized resources
2017-11-25 23:17:52,883 DEBUG org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpencilPlugin: Check: /tmp/hadoop-yarn/nm-local-dir/usercache/yarn/a
ppache/application_1511675189604_0010/filecache/12/hello_world.aocx
2017-11-25 23:17:52,883 DEBUG org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpencilPlugin: Found: /tmp/hadoop-yarn/nm-local-dir/usercache/yarn/
appache/application_1511675189604_0010/filecache/12/hello_world.aocx
2017-11-25 23:17:52,883 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.fpga.FpgaResourceHandlerImpl: IP file path: /tmp/hadoop-yarn/nm-local-dir/usercac
he/yarn/appache/application_1511675189604_0010/filecache/12/hello_world.aocx
2017-11-25 23:17:53,407 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.monitor.ContainersMonitorImpl: Skipping monitoring container container_1511675189604_0010_01_000
001 since CPU usage is not yet available.
2017-11-25 23:17:55,085 DEBUG org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpencilPlugin: aocl program: Running program from /home/fpga/intelF
PGA_pro/17.0/hld/board/nalla_pcie/linux64/libexec
Programming device: p385a_sch_axil15 : nalla_pcie (acInalla_pcie0)
Program succeed.

2017-11-25 23:17:55,085 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpencilPlugin: Intel aocl program /tmp/hadoop-yarn/nm-local-dir/user
cache/yarn/appache/application_1511675189604_0010/filecache/12/hello_world.aocx to aocl successfully
2017-11-25 23:17:55,085 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.fpga.FpgaResourceAllocator: Update IPID to hello_world for this allocated device
:FPGA Device:(Type: IntelOpenCL, Major: 247, Minor: 0, IPID: hello_world)

```

3. Test FPGA demo application with FPGA but no IP env specified

The YARN configuration is same with case 1.

Test command	Expectation	Actual Result
test_distributedshell_with _FPGA_without_IP.sh (-container_fpga_count 1)	FPGA allocation log	
	No FPGA IP download log	
	No FPGA IP programming log	
	Demo pass	

```





2017-11-25 22:47:20,691 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.fpga.FpgaResourceHandlerImpl: FpgaAllocation:
FpgaAllocation
  Allowed:
    FPGA Device:(Type: IntelOpenCL, Major: 247, Minor: 0, IPID: null)
  Denied

2017-11-25 22:47:20,693 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpencilPlugin: Got environment: , search IP file in localized resour
ces
2017-11-25 22:47:20,693 WARN org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpencilPlugin: IP ID environment is empty, skip downloading
2017-11-25 22:47:20,693 WARN org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.fpga.FpgaResourceHandlerImpl: FPGA plugin failed to download IP but continue, pl
ease set the value of environment variable: REQUESTED_FPGA_IP_ID if you want yarn to help

```

4. Test demo application with FPGA but no IP file uploaded

The yarn configuration is same with case 1.





Test command	Expectation	Actual Result
test_distributedshell_with_FPGA_without_IP.sh (-container_fpga_count 1 but remove -aocx option)	FPGA allocation log	
	No FPGA IP download log	
	No FPGA IP programming log	
	Demo pass(because it has devices and find aocx file in its directory)	

```

2017-11-26 00:16:19,569 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.fpga.FpgaResourceHandlerImpl: FpgaAllocation:
FpgaAllocation
  Allowed:
    FPGA Device:(Type: IntelOpenCL, Major: 247, Minor: 0, IPID: matrix_mult)
  Denied

2017-11-26 00:16:19,570 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpenclPlugin: Got environment: matrix_mult, search IP file in localized resources
2017-11-26 00:16:19,570 WARN org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.fpga.FpgaResourceHandlerImpl: FPGA plugin failed to download IP but continue, please set the value of environment variable: REQUESTED_FPGA_IP_ID if you want yarn to help

```

Test command	Expectation	Actual Result
test_distributedshell_with_FPGA_without_IP.sh (-container_fpga_count 1 but remove -aocx option and the environment setting)	FPGA allocation log	
	No FPGA IP download log	
	No FPGA IP programming log	
	Demo pass(because it has devices and find aocx file in its directory)	

```


2017-11-26 00:24:11,169 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.fpga.FpgaResourceHandlerImpl: FpgaAllocation:
FpgaAllocation
  Allowed:
    FPGA Device:(Type: IntelOpenCL, Major: 247, Minor: 0, IPID: matrix_mult)
  Denied





2017-11-26 00:24:11,170 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpenclPlugin: Got environment: , search IP file in localized resources
2017-11-26 00:24:11,170 WARN org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpenclPlugin: IP_ID environment is empty, skip downloading
2017-11-26 00:24:11,170 WARN org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.fpga.FpgaResourceHandlerImpl: FPGA plugin failed to download IP but continue, please set the value of environment variable: REQUESTED_FPGA_IP_ID if you want yarn to help

```

5. Test demo application without FPGA devices requested

The YARN configuration is same with case 1.

Test command	Expectation	Actual Result
test_distributedshell_with_out_FPGA_with_IP.sh	FPGA allocation log	

And test_distributedshell_with out_FPGA_without_IP.sh (-container_fpga_count 0)	No FPGA IP download log	
	No FPGA IP programming log	
	Demo failed (No device found)	
test_distributedshell.sh	"Env" command output	

```

2017-11-25 23:58:25,716 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.fpga.FpgaResourceHandlerImpl: container_1511675189604_0011_01_000002 requested 0
Intel FPGA(s)
2017-11-25 23:58:25,716 INFO org.apache.hadoop.yarn.server.nodemanager.containermanager.linux.resources.fpga.FpgaResourceHandlerImpl: FpgaAllocation:
FpgaAllocation
  Allowed:
  Denied:
FPGA Device:(Type: IntelOpenCL, Major: 247, Minor: 0, IPID: hello.world)

```

```

Container: container_1511675189604_0011_01_000002 on fpga-devkit.sh.intel.com_47676
LogAggregationType: AGGREGATED
=====
LogType:stdout
LogLastModifiedTime:Sat Nov 25 23:58:28 -0700 2017
LogLength:210
LogContents:
FPGA test:\n
Matrix sizes:
  A: 2048 x 1024
  B: 1024 x 1024
  C: 2048 x 1024
Initializing OpenCL
ERROR: CL_DEVICE_NOT_FOUND
Location: ../common/src/AOCLUtils/opencl.cpp:356
Query for number of devices failed
=====
End of LogType:stdout
*****








```



6. No fpga resource plugin configuration

```
<!--
<property>
  <description>
    Enable additional discovery/isolation of resources on the NodeManager,
    split by comma. By default, this is empty.
    Acceptable values: { "yarn-io/gpu", "yarn-io/fpga"}.
  </description>
  <name>yarn.nodemanager.resource-plugins</name>
  <value>yarn.io/fpga</value>
</property>
-->
<property>
  <description>
    When yarn.nodemanager.resource.fpga.allowed-fpga-devices=auto specified,
    YARN NodeManager needs to run FPGA discovery binary (now only support
    Intel's aocl) to get FPGA information.
    When value is empty (default), YARN NodeManager will try to locate
    discovery executable from system environment "ALTERAOCLSDKROOT"
  </description>
  <name>yarn.nodemanager.resource-plugins.fpga.path-to-discovery-executables</name>
  <value>/home/fpga/intelFPGA_pro/17.0/hld/bin/aocl</value>
</property>
<property>
  <description>
    Specify FPGA devices which can be managed by YARN NodeManager, split by comma
    Number of FPGA devices will be reported to RM to make scheduling decisions.
    Set to auto (default) let YARN automatically discover FPGA resource from
    system.

    Manually specify FPGA devices if admin only want subset of FPGA devices managed by YARN.
    At present, since we can only configure one major number in c-e.cfg, FPGA device is
    identified by their minor device number. A common approach to get minor
    device number of FPGA is using "aocl diagnose" and check uevent with device name.
  </description>
  <name>yarn.nodemanager.resource-plugins.fpga.allowed-fpga-devices</name>
  <value>auto</value>
</property>
<property>
  <name>yarn.nodemanager.resource-plugins.fpga.vendor-plugin.class</name>
  <value>org.apache.hadoop.yarn.server.nodemanager.containermanager.resourceplugin.fpga.IntelFpgaOpencilPlugin</value>
</property>
</configuration>
```

No fpga resource enabled configuration, it should run all scripts success because no FPGA isolation involved.

Test command	Expectation	Actual Result
start-yarn.sh	No error of NM	
test_distributedshell.sh	See output of "env"	
test_distributedshell_with_FPGA_with_IP.sh (-container_fpga_count 0)	A.No FPGA allocation log	
	B. No FPGA IP download log	
	C. NO FPGA IP programming log	
	D. Demo pass	
test_distributedshell_with_FPGA_without_IP.sh	Above "A, B, C, D" check items	

test_distributedshell_w ithout_FPGA_without_I P.sh	Above "A, B, C, D" check items	
test_distributedshell_w ithout_FPGA_with_IP. sh	Above "A, B, C, D" check items	

The demo binary output in container:

```
Container: container_1511624903085_0001_01_000002 on fpga-devkit.sh.intel.com_49726
LogAggregationType: AGGREGATED
=====
LogType:stdout
LogLastModifiedTime:Sat Nov 25 08:49:02 -0700 2017
LogLength:446
LogContents:
FPGA test:\n
Matrix sizes:
  A: 2048 x 1024
  B: 1024 x 1024
  C: 2048 x 1024
Initializing OpenCL
Platform: Intel(R) FPGA SDK for OpenCL(TM)
Using 1 device(s)
  p385a_sch_axl15 : nalla_pcie (aclnalla_pcie0)
Using AOCX: matrix_mult.aocx
Generating input matrices
Launching for device 0 (global size: 1024, 2048)

Time: 36.959 ms
Kernel time (device 0): 36.914 ms

Throughput: 116.21 GFLOPS

Computing reference output
Verifying
Verification: PASS

End of LogType:stdout
*****
```

Untested End-to-End Scenarios

At present, we haven't tested below scenarios. The report will be posted later.

1. Multiple devices in one host
2. A spark application that use native libraries to accelerate workload