Preliminary Performance Evaluation of Near-Data Processing in Caerus

This document summarizes the preliminary performance results of the near-data processing (NDP) capabilities in Project Caerus. The experiments measured the amount of network I/O and application latency when running the TPCH benchmark on a Spark.clutser, with data residing on a separate HDFS cluster. To enable NDP, we introduced our own DataSource module in Spark to plan for and perform computation pushdown. We also made slight modifications to the HDFS protocol to allow NDP-related information to be passed to the HDFS data nodes. The study compared three protocols: standard WebHDFS, modified HDFS without actual pushdown, and modified HDFS with HDFS pushdown.

The input dataset is 100GB in size, generated by ./dbgen -s 100. All tests use 4 Spark workers.

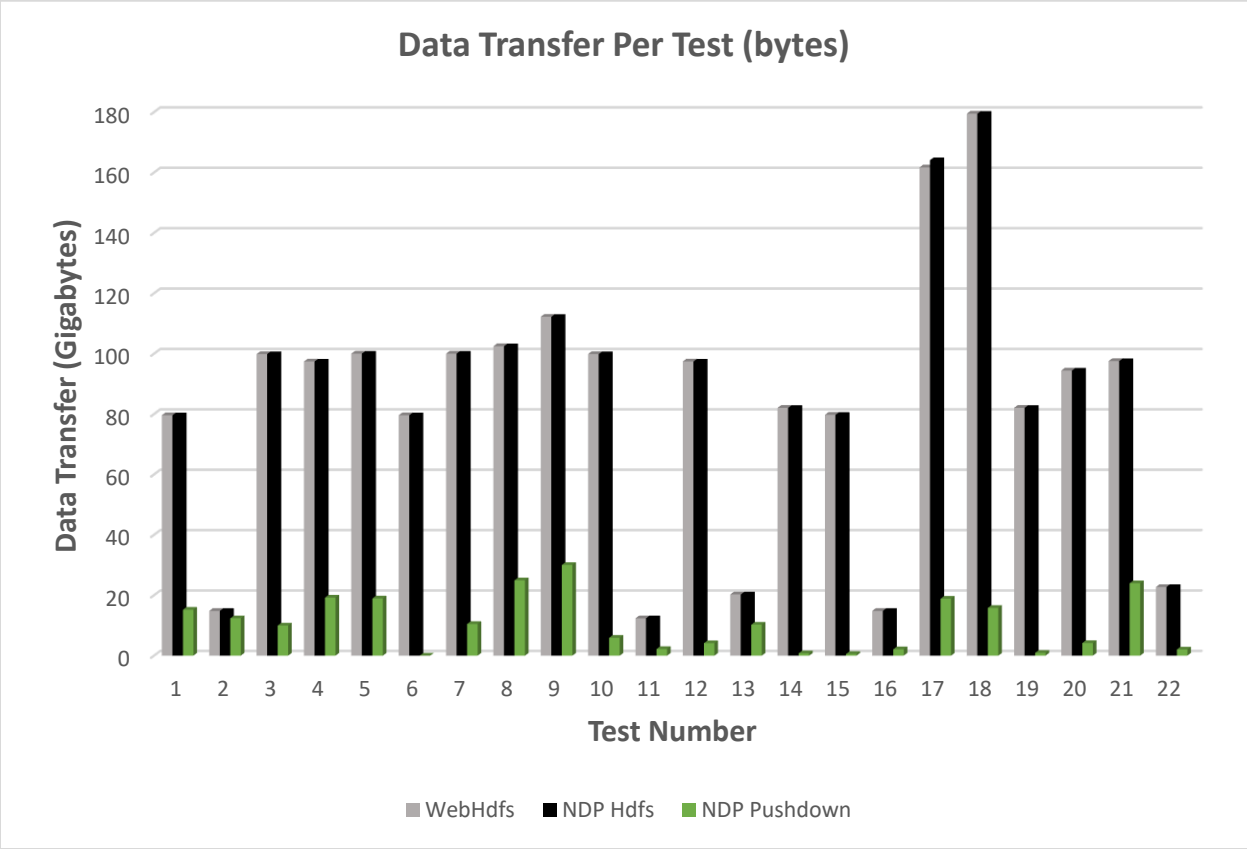
Data Transfer Per Test (bytes)
Data Transfer (Gigabytes)
180
160
140
120
100
80
60
40
20
0
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22
Test Number
WebHdfs
NDP Hdfs
NDP Pushdown

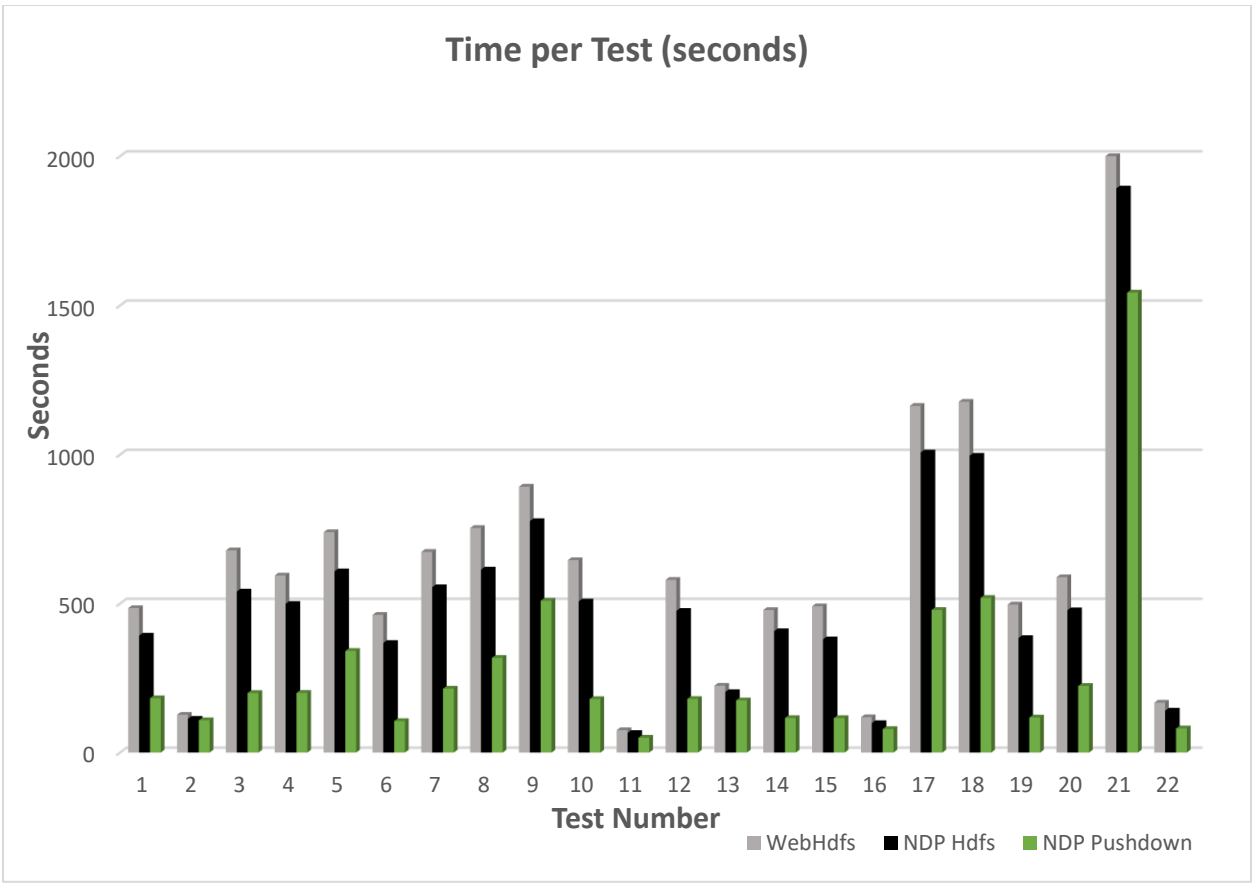
Time per Test (seconds)
2000
1500
1000
500
0
Seconds
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
Test Number
WebHdfs
NDP Hdfs
NDP Pushdown