




Comparing WANdisco Fusion to Cloudera BDR A Methodology and Results

November 10, 2015

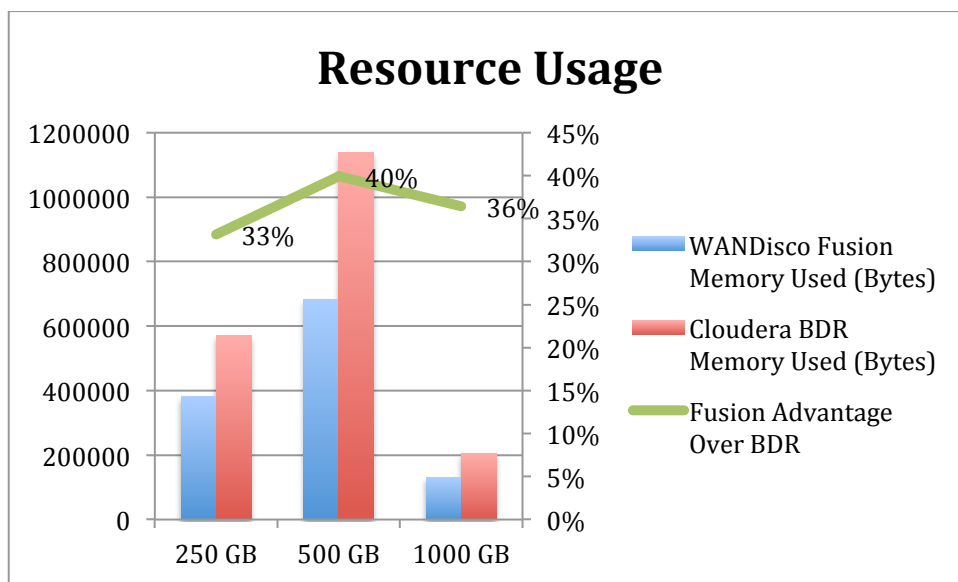
Executive Summary

Appvance implemented its PushToTest methodology to determine functional and performance differences between WANdisco Fusion and Cloudera BDR products for Hadoop. Appvance observes the following:

- 1) WANdisco Fusion provides functionality and availability not possible with Cloudera BDR. WANdisco's replication technology in Fusion enables clusters to be available even when backups are engaged, keeps clusters from losing data or diverging from other clusters over time, and provides continuous hot backup without downtime.
- 2) WANdisco provides faster backup speeds than Cloudera BDR. For example, Fusion takes 37% less time to backup 250 GB of data into a non-replicated path.



- 3) WANdisco uses less memory resources than Cloudera BDR. For example, Fusion takes 33% less memory to copy 250 GB of data.



The research and methodology presented in this paper enables enterprise architects, storage experts, development managers, and technical business/project managers to perform their own performance and functional tests.

Appvance offers the Appvance Performance Cloud (APC) and global testing services to Hadoop users, including WANdisco and Cloudera users. These tools and services deliver functional and performance metrics for an organization's business cost benefit analysis (CBA) projects. The results and supporting materials from this paper are found at: <http://appvance.com/hadoop>

When your organization chooses Hadoop then Appvance recommends WANdisco Fusion over Cloudera BDR for backup/restore and disaster recovery.

Background

Over the past few years distributed data caches and cluster-based distributed data environments became popular with large scale IT systems. Hadoop now provides the data scalability solution to Facebook and many other large systems. Hadoop provides distributed processing of large data sets across clusters of computers using simple programming models. Apache designed Hadoop to scale up from single servers to thousands of machines, each offering local computation and storage. Rather than rely on hardware to deliver high-availability, Hadoop software detects and handles failures at the application layer, to deliver a highly available data service on top of a cluster of computers, each of which may be prone to failures.

Methodology

This comparison follows the PushToTest methodology developed by Frank Cohen, founder at Appvance. The methodology breaks the comparison into two phases: Functional testing and Performance testing.

We accomplish functional tests by creating a test environment in a 2 cluster environment with 2 servers per cluster. We use bare metal dedicated servers hosted by IBM SoftLayer, with the following specifications:

- Dual Intel Xeon E5-2620 v3 (6 Cores, 2.40 GHz) (San Jose 1)
- Operating System: Windows Server 2012 Standard Edition (64 bit)

- RAM: 32 GB RAM
- Disk Controller: RAID
- First Hard Drive: 500 GB SATA II
- Second Hard Drive: 500 GB SATA II
- Third Hard Drive: 1.00 TB SATA II
- Fourth Hard Drive: 1.00 TB SATA II
- Fifth Hard Drive: 1.00 TB SATA II
- Public Bandwidth: 20000 GB Bandwidth
- Uplink Port Speeds: 1 Gbps Redundant Public & Private Network Uplink

On each server we installed WANdisco Fusion 2.6 and Cloudera BDR 5.4.5 binaries for CDH 5.4.5 / Ubuntu 14.04 LTS. Additionally, we installed Appvance Performance Cloud 2.2.174.

Appvance designed Appvance Performance Cloud (APC) in a console/controller architecture. For this test Appvance runs one console and 4 controllers. For the functional test the controllers operate command-line operations. For performance tests the controllers operate the Test Suite by calling Apache Teragen, Terasoft, TeraValidate.

The functional test suite performs the following tasks and checks for pass/fail conditions after each task:

1. Graceful shutdown of HDFS in a primary data center with two data centers.
2. Unexpected reboot of HDFS in primary data center with two data centers.
3. Graceful shutdown of HDFS in a primary data center with two data centers.
4. Unexpected reboot of HDFS in primary data center with two data centers.
5. Loss of network connectivity between primary and secondary data centers.
6. Automated recovery of Fusion once a network is restored.
7. Write data into primary and secondary data centers at the same time.

We accomplished performance testing by running the test environment and clusters on Amazon EC2 instances. 4 EC2 instances per cluster - 4 clusters - total 16 EC2 instances. EC2 instance type is m4.2xlarge. These come with 2.4 GHz Intel Xeon® E5-2676 v3 (Haswell) processors with 8 CPUs and 32 G memory. Each instance uses attached storage of 2 T bytes.

Performance tests use a data synchronization transfer tool between Hadoop instances for a backup and disaster recovery use case. We create Test Scenarios and their accompanying Test Use Cases to automate the test running in Appvance Performance Cloud (APC.) The Test Scenario in APC runs Teragen, Terasort, TeraValidate at 3 levels of data:

Small	250 Gb of data
Medium	500 Gb of data
Large	1 Tb of data

For each of the sizes of data (small, medium, large) we performed these tasks:

1. Run distcp to copy that data from the primary to the secondary data center.
2. Time hadoop distcp hdfs://primary/non-replicated hdfs://secondary/non-replicated.
3. Record the amount of time it takes for the operation to complete.
4. Do steps 1-3 using Cloudera BDR
5. Do steps 1-3 using WANdisco Fusion

The test use case measures the time it takes to operate the 3 test levels and provide a report comparing times between Cloudera (CDH) BDR and WANdisco Fusion replication.


Functional Testing Results

- Backup clusters may be fully utilized even when backups are engaged
 - ⇒ Even when the WANdisco Fusion is replicating data, clusters are still available. Cluster memory utilization is 23% less for Fusion than BDR. And, the cluster is continuously available for other processes.
- Active-Active offers benefits over Active-Passive
 - ⇒ Active-Active makes the data continuously available to prevent loss of data and also in meeting the pivotal SLAs.
 - ⇒ Data is made available despite Data Center failures.
 - ⇒ Recovery time is faster in the event of Data Center failures when compared to traditional BDR.
 - ⇒ Recovery Point Objective (RPO) is very low when compared to traditional BDR.
 - ⇒ Ingest and replicate at the same time with Fusion, not available with BDR.
- Higher availability when using replication over backup/restore, including less risk of data loss
 - ⇒ Replication ensures critical data is continuously available compared to traditional BDR
 - ⇒ No SPOF (Single Point Of Failure)
 - ⇒ Flapping WAN intermittent network and hardware failures, Fusion picks-up where it dropped off, including automatic recovery
 - ⇒ Forward recovery not available in Cloudera BDR, because it is a MapReduce job, does not start where it left off
 - ⇒ BDR resource contention means other apps are not allowed to run during the backup, can not do continuous hot backup
- Replication ensures clusters do not diverge over time
 - ⇒ Any failures in the either of the Hadoop clusters causes the file systems to diverge

Performance Test Results

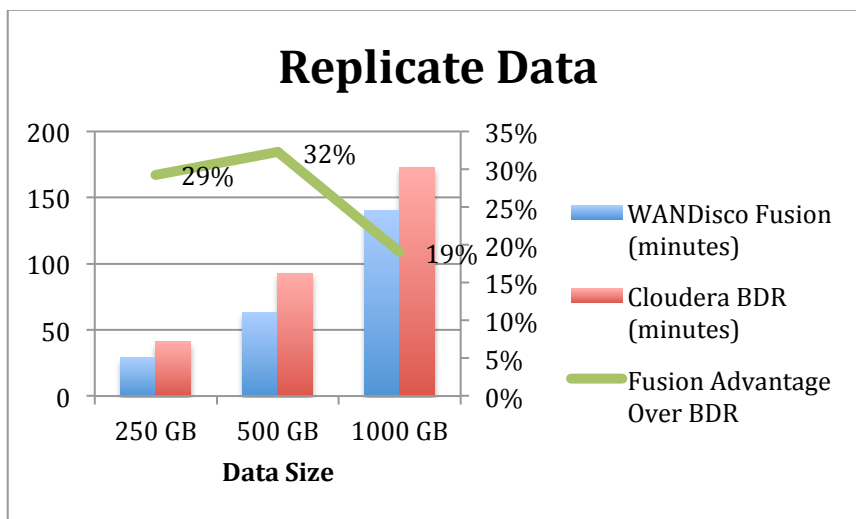
WANDisco provided faster backup speeds than Cloudera BDR. For example, Fusion takes 37% less time to backup 250 GB of data into a non-replicated path.

Perform A Backup	250 GB	500 GB	1000 GB
WANDisco Fusion (minutes)	72	99	110
Cloudera BDR (minutes)	114	125	141
Fusion Advantage Over BDR	37%	21%	22%




WANDisco Fusion using replication operated faster than Cloudera BDR as follows.

Replicate Data	250 GB	500 GB	1000 GB
WANDisco Fusion (minutes)	29	63	140
Cloudera BDR (minutes)	41	93	173
Fusion Advantage Over BDR	29%	32%	19%



WANdisco uses less resources (CPU, network, and memory) than Cloudera BDR. For example, Fusion takes 33% less memory to copy 250 GB of data.

Resource Usage	250 GB	500 GB	1000 GB
WANDisco Fusion Memory Used (Bytes)	382032	683381	131182
Cloudera BDR Memory Used (Bytes)	571371	1138173	206379
Fusion Advantage Over BDR	33%	40%	36%



How To Use The Results

Appvance offers the Appvance Performance Cloud (APC) and Global Test Services to Hadoop users, including WANdisco and Cloudera users. Appvance proves the functional and performance metrics for organizations doing cost benefit analysis (CBA) projects. The results and supporting materials are found at: <http://appvance.com/hadoop>

Appvance Global Services offers the methodology, tools, and experts used in this paper to you and your organization. For example, use the Appvance Performance Cloud system to test your Hadoop installations. Appvance offers the tests and APC environment under a Software As A Service (SAAS) and commercial software license, including training to your own team members.

Next Steps

As a follow-up report to this paper Appvance is considering extending the test to understand the scalability and performance in Apache Spark. Spark is a memory data cache that requires different memory models than a normal Hadoop instance. Appvance is considering a comparison of Spark and regular MapReduce jobs running slowly on a single cluster, then speeding-up when moved to separate clusters. Contact info@appvance.com to learn more about this.