Comparing WANdisco Fusion to Cloudera BDR
A Methodology and Results

April 2016

Executive Summary

Appvance implemented its PushToTest methodology to determine functional and performance differences between WANdisco Fusion and Cloudera BDR products for Hadoop. In summary, Appvance observed 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 full cluster backups are engaged, keeps clusters from losing data or diverging from other clusters over time so they remain consistent, and provides continuous hot backup and automated disaster recovery without downtime, even when entire data centers go offline. Fusion replication architecture avoids BDR's risk of data loss when outages occur after the last after-hours backup. For data access, Fusion offers higher data availability, especially critical in highly regulated industries where current and consistent backups of production data must be available at all times.

2) WANdisco Fusion's performance and resource utilization during replication are superior to Cloudera BDR's. Fusion performed a full cluster backup 37% faster than BDR, and used 40% less memory. During testing designed to simulate normal operation when other applications are running, WANdisco Fusion replicated the same data volumes up to 90% faster than BDR, without impacting the performance of the other applications.

Why Fusion Did So Well

Appvance attributes the performance, memory use, and availability advantages observed in Fusion to the following architectural and operational advantages:

1) WANdisco's active-active replication technology and its proxy server architecture performs replication on every change for user-selected HDFS folders. WANdisco replicates data in block and sub-block increments and does not require a file to be fully written and closed before replication. We anticipate streaming "Fast Data" applications where replication is required to benefit. BDR is built on MapReduce using a batch, script-driven architecture, requires files to be fully written and closed before data can be copied, and actively competes for cluster resources with other applications.

2) WANdisco Fusion provides operational advantages in five key areas:

   - **Deployment:** Fusion can be deployed on live production clusters without downtime. Fusion operates in-line as a proxy, running on an edge node as a client application to each Hadoop cluster, either on a physical server or a VM. Fusion requires no changes to the source code of the underlying Hadoop clusters, or other applications that run with them.

- **Data Replication:** Fusion detects every change and replicates the change automatically. Fusion provides full read/write access across any number of clusters and data centers. BDR requires scripts to be maintained and data movement is periodic and one-way. The same files cannot be written to everywhere, since there is no mechanism to maintain consistency across clusters if writes are allowed for the same files on all clusters.
- **Recovery:** Fusion provides automated DR. Clusters are resynchronized automatically, picking up where data replication left off after planned or unplanned outages. BDR requires manual intervention after an outage and may require the entire BDR job to be rerun, typically with checksums on, which consumes significant resources.
- **Network Administration:** Fusion only requires the Fusion servers to communicate through a firewall for remote replication to other data centers. BDR requires every data node in every cluster to be able to communicate with every other, which means that ports have to be configured for every single data node in every participating cluster if there is a firewall between them. This means that BDR imposes a significant burden on network security administrators as cluster size grows, and increases the potential attack surface open to hackers.
- **Flexibility:** Cloudera BDR copies data between specific versions of CDH clusters. WANdisco Fusion replicates data between Hadoop clusters running on any distribution, version and storage, including cloud storage, for systems supporting Hadoop Compatible File System (HCFS) API. Fusion avoids Hadoop distribution vendor lock-in, and enables source and target clusters to operate in parallel during migration, without production downtime.

Who Should Use This Report

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 data replication, backup 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 2.6.0 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, Terasort, 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 Test Results

1. 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 up to 40% less for Fusion than BDR. In addition, the cluster is continuously available for other processes.

2. 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 Service Level Agreements (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.

3. 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, cannot do continuous hot backup

4. Replication ensures clusters do not diverge over time

⇒ Any failures in the either of the Hadoop clusters cause the file systems to diverge with Cloudera BDR and not with WANdisco Fusion.

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%

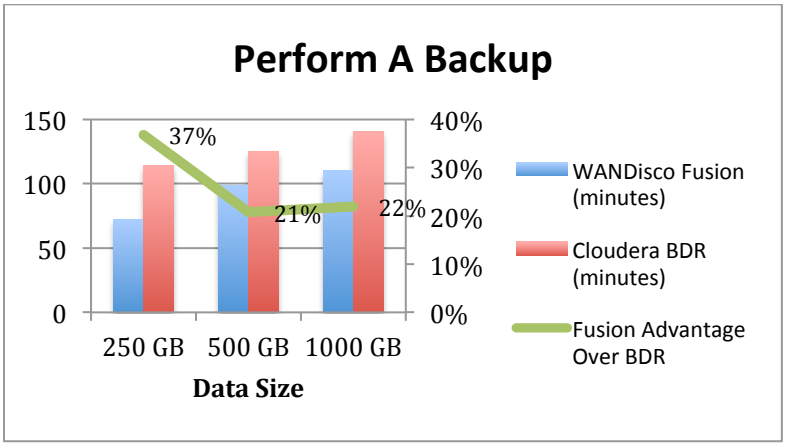

WANdisco uses less resource (CPU, network, and memory) than Cloudera BDR. For example, Fusion takes 40% less memory to copy 500 GB of data during a backup.

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%

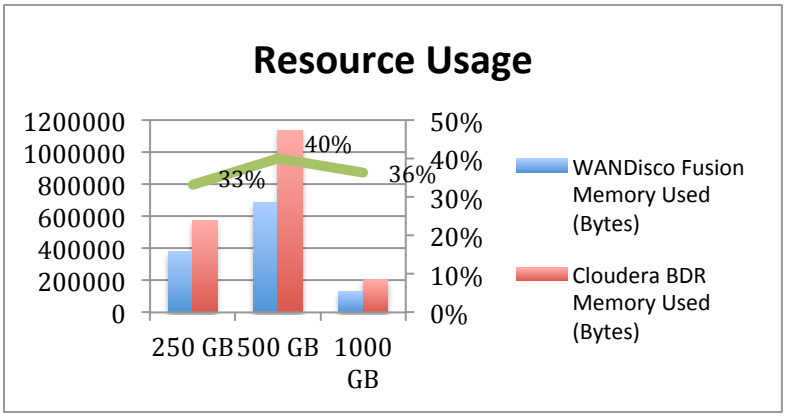

WANdisco transfers data up to 90% faster than BDR when other applications are running on the source cluster. BDR and Fusion were tested separately with the same ten MapReduce applications running on the source cluster for the data volumes shown below. This was done to simulate a typical Hadoop environment during normal operation.

Data volume	10 GB	25 GB	100 GB
WANdisco Fusion time (mm:ss)	00:29	00:44	4:06
Cloudera BDR time (mm:ss)	3:54	7:15	23:35
Fusion Advantage Over BDR	88%	90%	83%

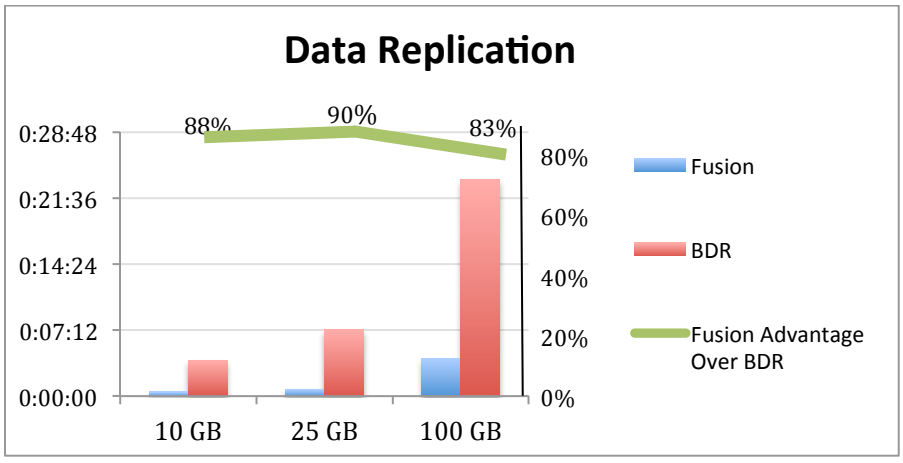

How to Use the Results

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 with “Fast Data” streaming applications such as those using 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.