

# A SURVEY ON CLOUD ANALYTICS AND THE FRAMEWORK TO SERVICE CLOUD ANALYTICS TECHNOLOGIES

VENKATESHWARA REDDY MUDIYALA  
*Master of Computer Science,  
 Department of CS, JNTU, Hyderabad, India*

**Abstract:** Cloud computing offers a powerful abstraction that provides a scalable, virtualized infrastructure as a service where the complexity of fine-grained resource management is hidden from the end-user. Running data analytics applications in the cloud on extremely large data sets is gaining traction as the underlying infrastructure can meet the extreme demands of scalability. Typically, these applications (e.g., business intelligence, surveillance video searches) leverage the MapReduce framework that can decompose a large computation into a set of smaller parallelizable computations. More often than not the underlying storage architecture for running a MapReduce application is based on an Internet-scale filesystem, such as GFS, which does not provide a standard (POSIX) interface. In this paper we revisit the debate on the need of a new non-POSIX storage stack for cloud analytics and argue, based on an initial evaluation, that it can be built on traditional POSIX-based cluster filesystems. In the course of the evaluation, we compare the performance of a traditional cluster file system and a specialized Internet file system for a variety of workloads for both traditional and MapReduce-based applications. We present modifications to the cluster filesystem's allocation and layout information to better support the requirements of data locality for analytics applications. We introduce the concept of a metablock that can enable the choice of a larger block granularity for MapReduce applications to coexist with a smaller block granularity required for traditional applications. We show that a cluster file system enhanced with metablocks can not only match the performance of specialized Internet file systems for MapReduce applications but also outperform them for traditional applications.

**Keywords :** Function shipping, MapReduce, Data analytics

## I. INTRODUCTION

Cloud computing is a compelling new paradigm that provides a scalable, virtualized infrastructure as a service, thereby enabling the end-user to exploit supercomputing power on demand without investing in huge infrastructure and management costs. This potential for unlimited scaling has made possible a plethora of cloud-based data analytics applications that can process extremely large sets of data.

These include newer applications for business intelligence, semantic web searches, video surveillance search, medical image analysis along with traditional data-intensive scientific applications such as satellite image pattern matching. A common feature in all these applications is that they are extremely parallel and their data access bandwidth requirements dominates other resource requirements. Such data-intensive applications where the computation can be easily decomposed into smaller parallel computations over a partitioned data set are a perfect match for Google's MapReduce framework [5] that provides a simple programming model using map and reduce functions over key/value pairs that can be parallelized and executed on a large cluster of machines. More recently, an open source version of MapReduce developed under the Apache Hadoop project is becoming a popular platform for building cloud data analytics applications. The underlying architecture for cloud computing typically comprises of large distributed clusters of low-cost servers in concert with a server virtualization layer and parallel programming libraries. One of the key infrastructure elements of the cloud stack, for data analytics applications, is a storage layer designed to support the following features: (1) scalable – to store petabytes of data, (2) highly reliable – to handle frequently-occurring failures in large systems, (3) low-cost – to maintain the economics of cloud computing, and (4) efficient – to best utilize the compute, network and disk resources.

The prevailing trend is to build the storage layer using an Internet scale filesystem such as Google's GFS [6] and its numerous clones including HDFS [1] and Kosmix's KFS [2]. The essential aspect of these filesystems is that they provide extreme scalability with reliability by striping and replicating the data in large chunks across the locally attached storage of the cluster servers, but simplify design and implementation by not providing a POSIX interface or consistency semantics.

Thus, they work well for MapReduce applications but cannot support traditional applications. We refer to such MapReduce focused file systems as specialized in the rest of the paper. In this paper, we revisit the debate on the need of a new non-POSIX storage stack for cloud analytics and argue, based on

an initial evaluation, that it can be built on traditional POSIX-based cluster filesystems. Existing deployments of cluster file systems such as Lustre [7], PVFS [3], and GPFS [8] show us that they can be extremely scalable without being extremely expensive. Commercial cluster file systems can scale to thousands of nodes while supporting 100 GBps sequential throughput. Furthermore, these file systems can be configured using commodity parts for lower costs without the need for specialized SANs or enterprise-class storage. More importantly, these file systems can support traditional applications that rely on POSIX file API's and provide a rich set of management tools. Since the cloud storage stack may be shared across different classes of applications it is prudent to rely on standard file interfaces and semantics that can also easily support MapReduce style applications instead of being locked in with a particular non-standard interface.

To this end, we address the challenges posed by the access characteristics of cloud analytics applications to traditional cluster file systems. First, we observe that MapReduce-based applications can co-locate computation with data, thus reducing network usage. We present modifications to the cluster filesystem's data allocation and data layout information to better support the requirements of data locality for analytics applications. Next, we observe that using large stripe unit sizes (or chunks) benefits MapReduce applications at the cost of other traditional workloads. To address that, we introduce a novel concept called metablock that can enable the choice of a larger block granularity for MapReduce applications to coexist with a smaller block granularity required for pre-fetching and disk accesses for traditional applications. While most analytics applications are read-intensive, we also enable write affinity that can better the performance of storing intermediate results by writing data locally. We compare the performance of both an Internet scale filesystem (Hadoop's HDFS) with a commercial cluster filesystem (IBM's GPFS) over a variety of workloads. We show that a suitably optimized cluster filesystem can match the performance of HDFS for a MapReduce workload (ideal data access pattern for HDFS) while outperforming it for the data access patterns of traditional applications. Concurrent to our work, researchers at CMU have undertaken an effort to provide support for Hadoop's MapReduce framework with PVFS [9]. It should be noted that we don't report HDFS performance for traditional file benchmarks since these benchmarks cannot be run on HDFS (even running with a FUSE layer only provides a subset of the POSIX interface).

## II. CHALLENGES

In this section, we evaluate the suitability of cluster file systems for cloud analytics applications. In our study, we selected for comparison the HDFS (Hadoop 18.1) filesystem

which is the de-facto filesystem for Apache's Hadoop project and IBM's GPFS cluster filesystem which is widely deployed in high-performance computing sites and whose source was readily available to us for modification. The hardware configuration we used is based on the IBM iDataPlex modular hardware architecture consisting of a single iDataPlex system with 42 nodes in two racks, where each node has 2 quad-core 2.2 GHz Intel Core2Duo CPUs, 8 GB RAM and 4 750 GB SATA drives. The nodes are connected by 2 Gigabit Ethernet switches (one per rack) with a 1 Gbps inter-switch link. The switch is Blade Network Technologies G8000 RackSwitch with 48 1 Gbps ports. The software running on each of these nodes is Linux 2.6.18 (CentOS 5.3) with two disks dedicated to the ext3 file system for storing intermediate results from computations and the remaining two disks dedicated to either GPFS or HDFS. We use 16 nodes in the experiment with 8 nodes on either rack.

### 2.1 Function shipping

The first drawback we found of cluster file systems is that they do not support shipping computation to data, a key feature exploited by the MapReduce class of applications [5]. In addition, the default block sizes are small which leads to a high task overhead for MapReduce applications that schedule one task per data block. To evaluate the effect of function shipping, we measured performance of a simple MapReduce grep application with GPFS and HDFS. The input to the grep application is a 16 GB text file. The Hadoop implementation did not take advantage of any block location information in GPFS and function shipping was not enabled as a result. Furthermore, we used the default block size of 64 MB in HDFS, whereas for GPFS we used a block size of 2 MB with pre-fetching turned on by default. The lack of co-location of computation with data, and the use of small blocks, are the main reasons for the slow-down in GPFS. The table shows that GPFS transfers several orders of magnitude more data over the network. In fact, the total amount of data transferred exceeds the input data size because of the default pre-fetching in GPFS. The filesystem sees 2 MB of data being read sequentially and pre-fetches multiple data blocks to satisfy expected reads. However, the map task for the next block may be scheduled on another node and thus most of the pre-fetched data is not used. High Availability. Another key requirement for data intensive applications is the ability to mask the failures of commodity components. Programs should be able to recover and progress in the event of multiple node and disk failures. This requires the data to be replicated across multiple nodes such that in the event of a node or disk failure, the computation can be restarted on a different node. Specialized file systems are designed based on this philosophy, and are able to tolerate multiple failures in the infrastructure. In comparison, cluster file systems have traditionally been

designed to use underlying data protection techniques (such as RAID) in shared storage to circumvent failures. However, the clusters that run data intensive applications typically do not use shared storage due to concerns regarding cost and bandwidth limitations, and instead attach local storage to each node in the cluster. While cluster file systems will run on nodes with locally attached storage (without replication or shared storage), the file systems will suffer data loss in the event of node or disk failures. Some cluster file systems (like GPFS) do provide data and meta-data replication as a means to survive node failures. The mechanism of replication can vary across file systems. GPFS, for example, uses a single source replication model, with the writer forwarding copies to all replicas. Specialized file systems, in contrast, use pipelined replication due to two important considerations: first, the out-bound bandwidth at the writer is not shared across multiple streams unlike the single-source model; second, write data can be pipelined in sequence from a node to the next node in the pipeline while the data is being written in the node.

For traditional applications, cluster file systems allow the use of concurrent writers to the same file, enabling the sharing of write bandwidth across multiple nodes. MapReduce applications usually have multiple simultaneous writers (to different files), so we don't expect the benefits of single-source replication to be significant. We hypothesize that it is possible for cluster file systems to match the write performance of specialized file systems and validate that in the experimental evaluation in Sections 4 and 5. However, we are continuing to explore the use of pipelined replication in cluster file systems

### III. METABLOCKS

Clearly, the grep application in the previous section demonstrated that running a MapReduce based application on a specialized file system has much better performance. In this section, we first attempt to mimic the basic properties of a specialized file system in GPFS and show the limitations of this approach. Next, we introduce the concept of a metablock, highlight the challenges in implementing the concept and demonstrate that GPFS is able to match the read performance of HDFS for MapReduce applications.

**Large blocks.** One approach would be to mimic the properties of specialized file systems as attempted in [9]. To achieve this, we increase the block size to a large value (16 MB) so that the map task and disk seek overhead is reduced (as one map task is assigned to each data block and will fetch the entire block for processing). Furthermore, we expose GPFS's block location information to the MapReduce layer in Hadoop so that tasks could be scheduled on the node where the data resides. In addition, we align the records in the input file with block boundaries, because a lack of alignment could result in

the fetch of a large data block just to read a partial record that straddles a block boundary. Finally, we turned pre-fetching off to avoid the network traffic of transporting large data blocks.

This particular version of GPFS is referred to as GPFS lb (GPFS with large blocks). To validate whether the approach would work, we use the same experimental setup as in Section 2 but with an input size of 80 GB. Figure 1 shows the relative performance of GPFS lb and HDFS in the experimental setup. The execution time of GPFS lb is almost the same as that of HDFS, but the network overheads of GPFS lb and HDFS are 2 GB and 1.4 GB of data transferred over the network during the duration of the experiment.

However, the performance parity with HDFS comes at a price. Turning off pre-fetching and making the unit of caching large in GPFS lb is detrimental to the performance of traditional filesystem workloads. Pre-fetching has been demonstrated to be extremely beneficial for sequential workloads and small block sizes are ideal for random workloads. To verify these effects, we compared unmodified GPFS to GPFS lb using the popular Bonnie filesystem benchmark [4]. The results of the experiment are shown in Table 2 and show a marked performance degradation for random workloads with the optimizations used in this section. There is an improvement for sequential workloads due to the large block size but the scale is not commensurate to the extent of the previously mentioned degradation. Bonnie also output other results that were consistent with the conclusions from the experiment.

**3.1 Metablocks:** The results of the evaluation indicate an interesting tradeoff in optimizing for data intensive and traditional applications. While a large block size is needed to minimize seek overheads and create a reasonable number of tasks in MapReduce applications, a small block size is needed for effective cache management and to reduce the pre-fetch overhead particularly when application records could span multiple blocks on different disks. Ideally, we need the best of both worlds where both seeks and pre-fetching are optimized so that both MapReduce and traditional applications can be supported. If the cluster file system could expose a large node-local block size to the MapReduce application and use a smaller block size for internal book-keeping, data transfer and pre-fetching, we can achieve the tradeoff. To better understand how we can manage this, we first describe the block allocation strategy used by GPFS.

GPFS implements wide-striping across the file system where large files are divided into equal sized blocks, and consecutive blocks are placed on different disks in a round-robin fashion. An allocation map keeps track of all disk blocks in the file system. To enable parallel updates to the allocation bit map, the map is divided into a large number of

lock-able allocation regions, with at least  $n$  regions for an  $n$  node system.

Each region contains the allocation status of  $l = n$ th of the disk blocks on every disk in the file system. This bitmap layout allows GPFS to allocate disk space properly striped across all disks by accessing only a single allocation region at a time. This approach minimizes lock conflicts because different nodes can allocate space from different regions. The allocation manager is responsible for keeping the free disk space statistics loosely- up-to-date across the cluster. To balance the block size selection tradeoff, we define a new logical construct called a metablock. A metablock is basically a consecutive set of blocks of a file that are allocated on the same disk. For example, 64 blocks of size 1 MB could be grouped into a 64 MB metablock. The GPFS round-robin block allocation is modified to use a metablock as the allocation granularity for the striping across the disks. Consequently, the block location map returned to the MapReduce application is also at the metablock granularity with the guarantee that all blocks in the metablock are in the same disk. Internally for all other pre-fetching and accesses, GPFS uses the normal block size granularity (which is 1 MB in our example). However there are two important challenges in implementing metablocks in GPFS – contiguity and fragmentation. First, it may not be possible to get a region with a set of blocks that is able to satisfy the contiguity requirement of a metablock.

In such a situation, the node trying to allocate a metablock will need to request a region with a contiguous set of blocks that can be used to build a metablock. However, a request to the allocation manager may incur network latency and affect the performance of a MapReduce application. To remedy the situation, a node prefetches a pool of contiguous regions ahead of time and requests new regions when the cardinality of the pool drops below a threshold. This means that a node will always have a ready pool of contiguous regions and will not incur network latency in the path of an I/O request. Second, allocating contiguous sets of blocks can lead to fragmentation in the allocation map, with skews in the free space of the regions in the allocation map depending on the nature of requests to the allocation manager. To address this issue, we relax the requirement for contiguity and only allocate sets of blocks which are contiguous only around 1 MB. This level of contiguity is reasonable for maintaining optimal sequential read and write performance, provides node-level locality, and minimizes the fragmentation problem. We evaluate the effectiveness of the metablock allocation scheme for the MapReduce grep application. The experimental setup and the input data size (80 GB) is identical to that in the previous experiment. Here, we experiment with HDFS (as

described before) and GPFS enhanced with a 16 MB metablock size and 512 KB block size, and no prefetching.

The results demonstrate that the execution time of GPFS with metablocks (referred to as GPFS mb) is within 10% of that of HDFS, while the network traffic is 2 worse than that of HDFS. Further, when we enabled controlled prefetching in GPFS (by specifying prefetch percentage), we incurred additional network traffic proportional to prefetch percentage.

A possible cause for concern is that the metablock optimization, which changes GPFS's allocation scheme, could have affected the performance of traditional applications. To confirm this hypothesis, we compared unmodified GPFS to GPFS mb. The results of the experiment are shown in Table 2 and show no marked difference between the two file systems.

The other results from Bonnie were also consistent with this result. Consequently, we conclude that metablocks do not hurt the performance of GPFS for traditional applications. It is important to note that this change to the allocation policy of the cluster file system does not impact the interface to the applications, and preserves the POSIX semantics provided by the unmodified system.

#### IV. MAPREDUCE FRAMEWORK

MapReduce is a framework for processing parallelizable problems across large datasets using a large number of computers (nodes), collectively referred to as a cluster (if all nodes are on the same local network and use similar hardware) or a grid (if the nodes are shared across geographically and administratively distributed systems, and use more heterogeneous hardware). Processing can occur on data stored either in a filesystem (unstructured) or in a database (structured). MapReduce can take advantage of the locality of data, processing it near the place it is stored in order to minimize communication overhead. A MapReduce framework (or system) is usually composed of three operations (or steps):

**Map:** each worker node applies the map function to the local data, and writes the output to a temporary storage. A master node ensures that only one copy of redundant input data is processed.

**Shuffle:** worker nodes redistribute data based on the output keys (produced by the map function), such that all data belonging to one key is located on the same worker node.

**Reduce:** worker nodes now process each group of output data, per key, in parallel.

MapReduce allows for distributed processing of the map and reduction operations. Maps can be performed in parallel, provided that each mapping operation is independent of the others; in practice, this is limited by the number of independent data sources and/or the number of CPUs near each source. Similarly, a set of 'reducers' can perform the reduction phase, provided that all outputs of the map operation that share the same key are presented to the same reducer at the same time, or that the reduction function is associative. While this process can often appear inefficient compared to algorithms that are more sequential (because multiple instances of the reduction process must be run), MapReduce can be applied to significantly larger datasets than a single "commodity" server can handle – a large server farm can use MapReduce to sort a petabyte of data in only a few hours.[13] The parallelism also offers some possibility of recovering from partial failure of servers or storage during the operation: if one mapper or reducer fails, the work can be rescheduled – assuming the input data are still available. Another way to look at MapReduce is as a 5-step parallel and distributed computation:

Prepare the Map() input – the "MapReduce system" designates Map processors, assigns the input key K1 that each processor would work on, and provides that processor with all the input data associated with that key. Run the user-provided Map() code – Map() is run exactly once for each K1 key, generating output organized by key K2. "Shuffle" the Map output to the Reduce processors – the MapReduce system designates Reduce processors, assigns the K2 key each processor should work on, and provides that processor with all the Map-generated data associated with that key. Run the user-provided Reduce() code – Reduce() is run exactly once for each K2 key produced by the Map step. Produce the final output – the MapReduce system collects all the Reduce output, and sorts it by K2 to produce the final outcome. These five steps can be logically thought of as running in sequence – each step starts only after the previous step is completed – although in practice they can be interleaved as long as the final result is not affected.

In many situations, the input data might already be distributed ("sharded") among many different servers, in which case step 1 could sometimes be greatly simplified by assigning Map servers that would process the locally present input data. Similarly, step 3 could sometimes be sped up by assigning Reduce processors that are as close as possible to the Map-generated data they need to process.

## V. DATA ANALYTICS IN CLOUD COMPUTING

Businesses have long used data analytics to help direct their strategy to maximize profits. Ideally data analytics helps eliminate much of the guesswork involved in trying to

understand clients, instead systemically tracking data patterns to best construct business tactics and operations to minimize uncertainty. Not only does analytics determine what might attract new customers, often analytics recognizes existing patterns in data to help better serve existing customers, which is typically more cost effective than establishing new business. In an ever-changing business world subject to countless variants, analytics gives companies the edge in recognizing changing climates so they can take appropriate action to stay competitive. Alongside analytics, cloud computing is also helping make business more effective and the consolidation of both clouds and analytics could help businesses store, interpret, and process their big data to better meet their clients' needs. The best uses of data analytics in cloud computing are:

**5.1 Social Media:** A popular use for cloud data analytics is compounding and interpreting social media activity. Before cloud drives became practical, it was difficult processing activity across various social media sites, especially if the data was stored on different servers. Cloud drives allow for the simultaneous examination of social media site data so results can be quickly quantified and time and attention allocated accordingly.

**5.1 Tracking Products:** Long thought of as one of the kings of efficiency and forethought, it is no surprise Amazon.com uses data analytics on cloud drives to track products across their series warehouses and ship items anywhere as needed, regardless of items proximity to customers. Alongside Amazon's use of cloud drives and remote analysis, they are also a leader in big data analysis services thanks to their Redshift initiative. Redshift gives smaller organizations many of the same analysis tools and storage capabilities as Amazon and acts as an information warehouse, preventing smaller businesses from having to spend money on extensive hardware.

**5.3 Tracking Preference:** Over the last decade or so, Netflix has received a lot of attention for its DVD deliver service and the collection of movies hosted on their website. One of the highlights of their website is its movie recommendations, which tracks the movies users watch and recommends others they might enjoy, providing a service to clients while supporting the use of their product. All user information is remotely stored on cloud drives so users' preferences do not change from computer to computer. Because Netflix retained all their users' preferences and tastes in movies and television, they were able to create a television show that statistically appealed to a large portion of their audience based on their demonstrated taste. Thus in 2013, Netflix's House of Cards became the most successful internet-television series ever, all thanks to their data analysis and information stored on clouds.

## VI. CLOUD ANALYTICS SERVICES

We can make fast, accurate data based decisions with analytics services. IBM Cloud provides data and AI services and tools. It features embedded intelligence capabilities through machine learning. The services provided by the cloud are the cloud is used to collaborate in teams across functions to access all trusted data and best-in-class technologies. IBM cloud use multiple analytics technologies to learn from the data and quickly get new answers for the business. IBM cloud used to deliver new insights to the business quickly and continuously improve them through rapid iteration.

## VI. FUTURE OPTIMIZATIONS

The most important was trying to make writes as network efficient in GPFS as they are in HDFS (due to the first replica being written to the local node). We designed an extension to metablocks which has allowed GPFS to potentially match the performance of HDFS for writes as well. The extension involves adding an ioctl call to GPFS which lets an application specify the set of hosts to be used by the metablock allocation scheme for a particular file. This allows Hadoop applications to specify that the first copy of data should reside on the local host, which is the policy used by HDFS. This technique reduces the network traffic during writes, and significantly improves write performance (up to a factor of 5). True to our theme, we use GPFS with prefetching enabled to benefit traditional as well as MapReduce workloads. This, however, exposes two interesting questions we are currently exploring: (1) Can we design an adaptive prefetching scheme such that it only consumes spare network bandwidth, and does not contend with critical network traffic? (2) Can any MapReduce workloads benefit from such prefetching, thereby outperforming HDFS? Similarly, we are also pursuing use cases of MapReduce workloads where GPFS, can in fact, outperform HDFS by leveraging features unique to a true file system such as ability to cope with client-side caching, and simultaneously support random and sequential workloads.

## VII. CONCLUSION

This paper evaluates the debate whether cluster file systems can potentially match the performance of Internet scale filesystems for cloud-based analytics applications. We examine the requirements of data intensive applications and show that cluster file systems are deficient in support for large block sizes and exposing block location information to MapReduce applications. To remedy this, we introduce the concept of metablocks that provide the illusion of large blocks for MapReduce applications, while providing the benefits of small blocks for traditional applications at the same time. We

show that a cluster file system enhanced with metablocks can provide the best of both worlds performance.

## VIII. REFERENCES

- [1] Hadoop Distributed Filesystem. <http://hadoop.apache.org>.
- [2] Kosmos Filesystem. <http://kosmosfs.sourceforge.net/>.
- [3] Parallel Virtual Filesystem. <http://www.pvfs.org/>.
- [4] The Bonnie Filesystem Benchmark. <http://www.textuality.com/bonnie/>.
- [5] J. Dean and S. Ghemawat. Mapreduce: Simplified data processing on large clusters. In Sixth Symposium on Operating System Design and Implementation, pages 137–150, December 2004.
- [6] S. Ghemawat, H. Gobioff, and S. Leung. The Google file system. In ACM SOSP, October 2003., 2003.
- [7] Lustre. The lustre storage architecture. <http://www.lustre.org/>.
- [8] F. Schmuck and R. Haskin. GPFS: A shared-disk file system for large computing clusters. In Proc. of the First Conference on File and Storage Technologies (FAST), pages 231–244, Jan. 2002.
- [9] W. Tantisiriroj, S. Patil, and G. Gibson. The crossing the chasm: Sneaking a parallel file system into hadoop. In SC08 Petascale Data Storage Workshop, 2008.