



## VIVA-TECH INTERNATIONAL JOURNAL FOR RESEARCH AND INNOVATION

ANNUAL RESEARCH JOURNAL

ISSN(ONLINE): 2581-7280

---

### Google File System

Pramod A. Yadav<sup>1</sup> , Prof. Krutika Vartak<sup>2</sup>

<sup>1</sup>(Department of MCA, Viva School Of MCA/ University of Mumbai, India)

<sup>2</sup>(Department of MCA, Viva School Of MCA/ University of Mumbai, India)

---

**Abstract :** Google File System was innovatively created by Google engineers and it is ready for production in record time. The success of Google is attributed to the efficient search algorithm, and also to the underlying commodity hardware. As Google runs a number of applications, its goal became to build a vast storage network out of inexpensive commodity hardware. So Google created its own file system, named as Google File System, which is GFS. Google File System is one of the largest file systems in operation. Generally, Google File System is a scalable distributed file system for large distributed data-intensive applications. In the design phase of Google File System, the given stress includes component failures, files are huge, and files are mutated by appending data. The entire file system is organized hierarchically in directories and identified by pathnames. The architecture comprises of multiple chunk servers, multiple clients, and a single master. Files are divided into chunks, and that is the key design parameter. Google File System also uses leases and mutation order in its design to achieve atomicity and consistency. As for fault tolerance, Google File System is highly available, replicas of chunk servers, and a master exist.

**Keywords** - Design, Scalability, Fault tolerance, Performance, Measurements.

---

#### I. INTRODUCTION

Google File System (GFS) to meet the ever-increasing needs of Google data processing needs. GFS shares many similar goals with previously distributed file systems such as functionality, feasibility, reliability, and availability. However, its design is driven by the critical recognition of our application workloads and technology, current and hosted, reflecting the significant departure from other design of file system design. We also reviewed the major sector options and explored very different points in the design space.

First, partial failure is more common than different. The file system contains hundreds or thousands of storage devices built from the most inexpensive components of the commodity type and is available with an equal number of customer's equipment. The quantity and quality of computers is likely to ensure that some do not work at any time and some will not recover from their rent failure. We've identified problems caused by application bugs, operating system bugs, human errors, and disks, memory, connectors, connections, and power of supplies. Therefore, regular monitoring, error detection, error tolerance, and automatic recovery should be in line with the system.

Second, the files are large by traditional standards. Multi-GB files are standard. Each file typically contains multiple application elements such as web documents. When we are constantly working on fast-growing data sets many TBs contain billions of items, it is impossible to manage billions of files equivalent to KB almost even if the file system can support them. For this reason, design considerations and parameters such as I/O performance and block size should be reviewed.

Third, many files are converted by inserting new data instead of overwriting existing data. Random write inside a file is almost non-existent. Once written, files are read only, and usually only in sequence. Various details share these features. Some may create large repositories where data analysis systems test. Some may be data streams that continue continuously through applications. Some may be historical data. Some may be temporary effects produced on one machine and processed in another, either simultaneously or later. Given this pattern of access to large files, the application becomes the cause of performance and atomic verification, while data storage for customers loses appeal.

Many GFS collections are currently used for a variety of purposes. The largest have more than 1000 storage areas, more than 300 TB of disk storage, and are widely available by hundreds of clients on continuously different machines.

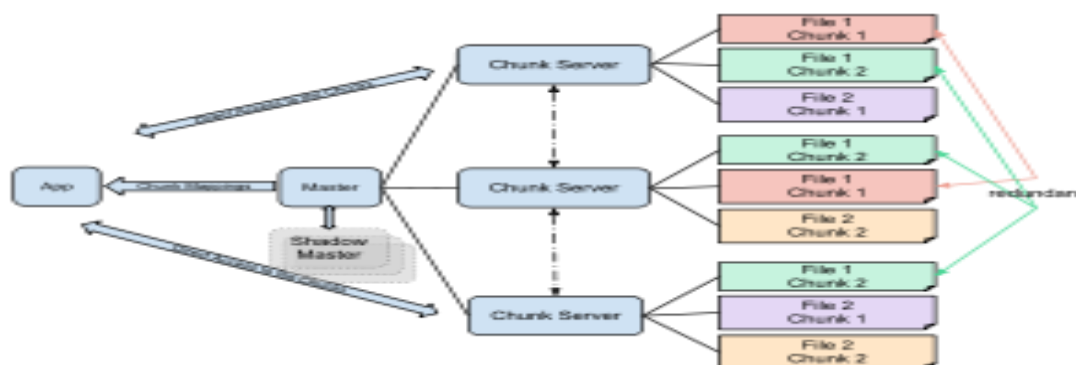


Figure1 : Google File System

## II. SYSTEM INTERACTIONS

### 2.1. Leases and Mutation Order

Conversion is a modification that changes chunk content or metadata such as writing or append-opion. Each modification is performed on all chunk pieces. We use the rental to maintain a consistent variable density in all replicas. The king offered a chunk rent to one of the repli-cas, which we call the principal. The primary is selected in serial sequence with all feature changes in the chunk. All drawings follow this order when inserting genetic modification. Therefore, the land reform order is defined first by the king's rental grant order, and within the lease serial numbers provided by the primary.

### 2.2. Data Flow

We cut the flow of data in the flow of control so we can use the network efficiently. While control flows from client to primary and then to all managers, data is pushed to the line in a carefully collected line of chunkservers in the form of pipelines. Our goals are to take full advantage of the network bandwidth of each machine, to avoid network barriers and high-latency links, and to reduce push-up delays in all data.

To take full advantage of the network bandwidth of each machine, data is pushed sequentially to the chunkservers line rather than distributed to another type of topology (e.g., Tree). Therefore, the full output bandwidth of each machine is used to transmit data much faster than is divided between multiple receivers.

To avoid network barriers and high-latency connections (e.g., inter-switch connectors are common to both) as much as possible, each machine transmits data to a machine that is "too close" to a network that has not yet received it. Our network topology is simple enough that "distances" can be accurately measured from IP addresses.

### 2.3. Atomic Record Appends

GFS provides atomic insertion known as record append In traditional writing, the client specifies the set in which the data will be recorded. Simultaneous writing in the same region is immeasurable: a region may end up containing pieces of data from multiple clients. In the appendend of the record, however, the client only

specifies the data.

Record recording is widely used by our distributed applications where multiple clients of different machines stick to the same file at the same time. Clients will need complex and expensive synchronization of synchronization, of getting more power using a locked manager, if they do so with traditional writing. In our workloads, such files are often it works as a multi-producer line / individual customer line or integrated results from various customers.

Record append is a form of transformation and follows the flow of control with just a little more sense in the beginning. The client pushes the data into all the final file fragments. After that, it sends its request to primary. The main test is to see if inserting a record in the current fragment could cause the chunk to exceed the maximum size (64 MB). If so, add a chunk to the maximum size, tell the seconds to do the same, and then respond to the client indicating that the task should be repeated in the next chunk exactly where he found it, and finally respond to the client's success.

#### **2.4.Snapshot**

The snapshot function makes a copy of the file or directory tree ("source") almost immediately, while mimicking any disruption of ongoing changes. Our users use it to quickly create branch copies of large data sets.

We use standard copying techniques in writing using abbreviations. When the master receives a summary request, it first removes any prominent leases in chunks from the files that are about to be downloaded. This ensures that any subsequent writing of these cases will require contacting the landlord to find a lease owner. This will give the king a chance to make a new copy of the chunk first.

After the lease is completed or expired, the master installs the disk functionality. It then uses this log to log in to its memory state by repeating the metadata of the source file or directory tree. Recent shooting files identify the same nodes as source files.

### **III. MASTER OPERATION**

#### **3.1.Namespace Management and Locking**

We now show you how this lock method can prevent file / home / user / foo from being created while / home / user is being downloaded to / save / user. The snapshot operation finds read lock in / home and / save, and locks in / home / user and / save / user. The structure of the acquires file is learned to lock in / home and / home / user, and to lock in / home / user / foo. The two activities will be designed to work well together as they try to find conflicting locks on the / home / user. File formatting does not require a text lock in the parent directory because there is no "directory", or inode-like, data structure that should be protected from conversion. The key read in the name is able to protect the parent's guide from removal.

#### **3.2.Replica placement**

The GFS collection is still widely distributed at more than one level. It usually has hundreds of chunkservers scattered throughout the machine racks. These chunkservers can also be accessed by hundreds of clients from the same or different racks. Connections between two machines in different markets may require one or more network switching. In addition to the ally, the bandwidth entering or exiting the rack may be less than the combined bandwidth of all equipment within the rack. Multi-level distribution presents a unique challenge in distributing data on robustness, reliability, and availability.

#### **3.3.Creation, Re-replication, Rebalancing**

Chunk duplication is created for three reasons: chunk duplication, duplication, and redesign.

When the king creates a lump, he chooses where you can put the empty answers first. It looks at several factors. (1) We want to place new duplicates on chunkservers with less than disk usage. Over time this will measure disk usage across all chunkservers. (2) We want to limit the number of "recent" items in each chunkserver. Although creation itself is cheap, it accurately predicts future heavy writing because episodes are made when they are sent in writing, and in append-once-read-most of our work is always read once they have read it in full. (3) As mentioned above, we want a partial distribution across all racks.

#### **3.4.Garbage Collection**

After a file is deleted, GFS does not immediately reclaim the available physical storage. It does so only lazily during regular garbage collection at both the file and chunk levels. We find that this approach makes the system much simpler and more reliable.

### **3.5.Mechanism**

When a file is deleted by the app, the master logs are removed as quickly as other changes. However, instead of retrieving resources quickly, the file is renamed a hidden name that includes a delete stamp. During the master scan of the space file name, it removes any hidden files if they have been lost for more than three days (time is corrected). Until then, the file can still be read under a new, special name and can be deleted by renaming it back to normal. When a hidden file is deleted from the namespace, its memory metadata is erased. This removes its links from all its links.

## **IV. FAULT TOLERANCE AND DIAGNOSIS**

### **4.1High Availability**

Among hundreds of servers in a GFS cluster, some are bound to be unavailable at any given time. We keep the overall system highly available with two simple yet effective strategies: fast recovery and replication

### **4.2.Fast Recovery**

Both king and chunkserver are designed to retrieve their status and start in seconds no matter how they cut. In fact, we do not distinguish between normal and abnormal cuts; servers are constantly shut down by the process kill. Clients and other servers experience a slight hiccup as they expire on their outstanding applications, reconnect with the restart server, and try again. Section 6.2.2 reports detected startup times.

### **4.3.Chunk Replication**

Users can specify different levels of duplication of different parts of file name space. Default three times. Master clones replicas are available as required to keep each cluster fully converted as chunkservers move. While repetition has worked well for us, we are exploring other forms of cross-server redundancy such as parity codes or erasure of our growing read-only retention needs. We expect it to be a challenge but it is manageable to use these sophisticated programs to redefine the performance of our freely integrated program because our campaign is dominated by appenders and learns rather than writing a little random.

### **4.4.Master Replication**

The main situation is honestly repeated. Work log and test locations are duplicated on multiple machines. A regime change is considered a commitment only after their entry record is placed on a local disk and in all major responses. For convenience, one major process always handles all modifications and background functions such as garbage collection that transforms the system internally. If it fails, it can restart almost immediately. If its machine or disk fails, monitoring infrastructure outside the GFS initiates a new process in another location with a duplicate operation. Clients use only the canonical name of the master (e.g. gfs-test), which is the alias of the DNS that can be changed when the master is transferred to another machine.

### **4.5.Diagnostic Tools**

The detailed and detailed login helped the im-equate with problem classification, error correction, and performance analysis, while only low cost is available. Outgoing logs, it is difficult to comprehend the passing, non-recurring connections between machines. GFS servers produce diagnostic logs that record many important events (such as up and down chunkservers) and all RPC requests and responses. These diagnostic logs can be removed freely without affecting the accuracy of the system. However, we try to keep these logs close to space permits.

## **V. MEASUREMENTS**

### **5.1.Micro-benchmarks**

We rated performance on a GFS collection consisting of one master, two master replicas, 16 chunkservers, and 16 clients. Note that this setting is set to facilitate testing. Standard collections have hundreds of chunkservers and hundreds of clients.

All devices are equipped with two 1.4 GHz PIII processors, 2 GB of memory, two discs of 80 GB 5400 rpm, and a 100 Mbps full-duplex Ethernet connection to the HP 2524 switch. All 19 GFS server equipment is connected. In one switch, and all 16 customer equipment to another. The two switches are connected via a 1 Gbps link.

### 5.2. Reads

Clients read simultaneously from the file system. Each client reads a randomly selected 4 MB region from a 320 GB file set. This is 256 times more so for each client to end up reading 1 GB of data. Chunkservers bundled together have only 32 GB of memory, so we expect at least 10% of Linux cache memory. Our results should be close to the cold storage results.

The combined reading rate of  $N$  clients and its theoretical limit. The limit is up to 125 MB / s when 1 Gbps link between these switches is full, or 12.5 MB / s per client when its 100 Mbps network interface is full, whichever works. The rated reading rate is 10 MB / s, or 80% of each client limit, where one client reads. The combined reading rate is up to 94 MB / s, approximately 75% of the link limit of 125 MB / s, for 16 students, or 6 MB / s per client. Technology decreases from 80% to 75% because as the number of students increases, so do the chances of more students learning simultaneously from the same chunkserver.

### 5.3. Writes

Clients write simultaneously in separate files of  $N$ . Each client writes 1 GB of data to a new file in a 1 MB write series. The level of writing combined with the theoretical limit is shown in Figure 3 (b). Fields are limited to 67 MB / s because we need to write each byte in 3 out of 16 chunkservers, each with an input connection of 12.5 MB / s.

One client's write rate is 6.3 MB / s, about half the limit. A major feature of this is our network stack. It does not work well with the piping system we use to compress data in chunk duplication. Delays in transmitting data from one to another reduce the overall writing rate.

### 5.4. Appends versus Writes

Appendices are widely used in our production programs. In the Cluster X, the average recording rate included is 108: 1 per relay by 8: 1 per performance count. In Group Y, which is used by production systems, the ratings are 3.7: 1 and 2.5: 1 respectively. Moreover, these ratios suggest that in both collections the appends tend to be larger than the writing. In Cluster X, however, the total usage of the recording time is relatively low so the results may be offset by one or two application settings by selecting the bus size.

## VI. CONCLUSIONS

We started by re-examining file system assumptions in terms of our current and anticipated operational and technological environment. Our vision has led to very different points in the construction space. We treat component failures as normal rather than different, prepare large files that are heavily loaded (perhaps simultaneously) and read (usually sequentially), and both expand and release the file system interface to improve the entire system.

Our system provides error tolerance by constantly scanning, duplicating important data, and fast and automatic updates. Chunk duplication allows us to tolerate chunkserver failure. The frequency of these failures encourages the online repair process of the novel which reevaluates the damage and transparency and compensates for lost answers very quickly. Additionally, we use data corruption testing tests at the disk or IDE subsystem level, which is most commonly assigned to the number of disks in the system.

Our design brings the full inclusion of many similar readers and writers who perform a variety of tasks. We do this by separating file system control, which passes through the master, from the data transfer, which passes directly between the chunkservers and customers. Master's involvement in normal operation is reduced by large chunk size and chunk rental, which authorizes primary duplication in data conversion. This makes it possible to be a simple, a masterized master not a bottle. We believe that the development of our network stack will remove the current limit on the transfer of writing seen by each client.

GFS has successfully met our storage needs and is widely used within Google as the ultimate platform for research and development and processing of production data. It is an important tool that enables us to continue to create and attack problems on the entire web scale.

### ACKNOWLEDGEMENTS

I thank my college for giving us the opportunity to make this project a success. I offer my special thanks and sincerity.

I thank Professor Krutika vartak for encouraging me to complete this research paper, for guidance and assistance for all the problems I encountered while doing research. Without his guidance, I would not have completed my research paper.

### REFERENCES

- [1] Thomas Anderson, Michael Dahlin, Jeanna Neefe, David Patterson, Drew Roselli, and Randolph Wang. Serverless network file systems. In Proceedings of the 15th ACM Symposium on Operating System Principles, pages 109–126, Copper Mountain Resort, Colorado, December 1995.
- [2] Remzi H. Arpaci-Dusseau, Eric Anderson, Noah Treuhaft, David E. Culler, Joseph M. Hellerstein, David Patterson, and Kathy Yelick. Cluster I/O with River: Making the fast case common. In Proceedings of the Sixth Workshop on Input/Output in Parallel and Distributed Systems (IOPADS '99), pages 10–22, Atlanta, Georgia, May 1999.
- [3] Luis-Felipe Cabrera and Darrell D. E. Long. Swift: Using distributed disk striping to provide high I/O data rates. *Computer Systems*, 4(4):405–436, 1991.
- [4] Garth A. Gibson, David F. Nagle, Khalil Amiri, Jeff Butler, Fay W. Chang, Howard Gobioff, Charles Hardin, Erik Riedel, David Rochberg, and Jim Zelenka. A cost-effective, high-bandwidth storage architecture. In Proceedings of the 8th Architectural Support for Programming Languages and Operating Systems, pages 92–103, San Jose, California, October 1998.
- [5] John Howard, Michael Kazar, Sherri Menees, David Nichols, Mahadev Satyanarayanan, Robert Sidebotham, and Michael West. Scale and performance in a distributed file system. *ACM Transactions on Computer Systems*, 6(1):51–81, February 1988.
- [6] InterMezzo. <http://www.inter-mezzo.org>, 2003.
- [7] Barbara Liskov, Sanjay Ghemawat, Robert Gruber, Paul Johnson, Liuba Shrira, and Michael Williams. Replication in the Harp file system. In 13th Symposium on Operating System Principles, pages 226–238, Pacific Grove, CA, October 1991.
- [8] Lustre. <http://www.lustreorg>, 2003.
- [9] David A. Patterson, Garth A. Gibson, and Randy H. Katz. A case for redundant arrays of inexpensive disks (RAID). In Proceedings of the 1988 ACM SIGMOD International Conference on Management of Data, pages 109–116, Chicago, Illinois, September 1988.
- [10] rank Schmuck and Roger Haskin. GPFS: A shared-disk file system for large computing clusters. In Proceedings of the First USENIX Conference on File and Storage Technologies, pages 231–244, Monterey, California, January 2002.
- [11] Steven R. Soltis, Thomas M. Ruwart, and Matthew T. O'Keefe. The Gobal File System. In Proceedings of the Fifth NASA Goddard Space Flight Center Conference on Mass Storage Systems and Technologies, College Park, Maryland, September 1996.
- [12] Chandramohan A. Thekkath, Timothy Mann, and Edward K. Lee. Frangipani: A scalable distributed file system. In Proceedings of the 16th ACM Symposium on Operating System Principles, pages 224–237, Saint-Malo, France, October 1997.