Big Data Analytics on Large Scale Shared Storage System

Kyar Nyo Aye
University of Computer Studies, Yangon, Myanmar
kyarnyoaye@gmail.com

Abstract

In today's world, almost every enterprise is seeing an explosion of data. They getting huge amount of digital data generated daily. Such huge amount of data needs to be stored for various reasons. Now the important question that arises at this point of time is how do we store, manage, process and analyze such huge amount of data most of which is Semi structured or Unstructured in a scalable, fault tolerant and efficient manner. The challenges of big data are most of them is semi structured or unstructured data, need to carry out complex computations over big data and the time required to process big data is as low as possible. In this paper, we propose big data platform based on Hadoop MapReduce framework and Gluster file system over large scale shared storage system to address these challenges. Our big data platform can support large scale data analysis efficiently and effectively.

Keywords: big data, big data analytics, big data platform, MapReduce framework, Gluster file system, scale-out NAS

1. Introduction

Today, information is generated continuously around the globe 24/7. Almost every growing organization wants to automate most of its business processes and is using IT to support every conceivable business function. This is resulting into huge amount of data being generated in the form of transactions and interactions. Web has become an important interface for interactions with suppliers and customers generating the huge amount of data in the form of emails etc. Besides this, there is a huge amount of data emitted automatically in the form of logs like network logs and web server logs.

Various Telecom Service Providers get huge amount of data in the form of conversations and Call Data Records. Various Social N/W Sites have started getting TBs of data every day in the form of tweets, blogs, comments, photos and videos etc. Facebook generates 4TBs of compressed data every day. Web Companies like these get huge amount of click stream data generated daily as well. Hospitals have data about the patients, their diseases and the data generated by various medical devices as well. Sensors used in various machines used for production keep generating so much of event data in seconds. Almost every sector like transport, finance is seeing a tsunami of data.

Such huge amount of data needs to be stored for various reasons. Sometimes any compliance demands more historical data to be stored. Sometimes organizations want to store, process and analyze this data for intelligent decision making to get the competitive advantage. For example analyzing CDR data can help a service provider know their quality of service and then make the necessary improvements. A credit card company can analyze the customer transactions for fraud detection. Server logs can be analyzed for fault detection. Web logs can help understand the user navigation patterns. Customer emails can help understand the customer behavior, interests and some time the problems with the products as well.

Now the important question that arises at this point of time is how do we store and process such huge amount of data most of which is Semi structured or Unstructured. There is a high-level categorization of big data platforms to store and process them in a scalable, fault tolerant and efficient manner [11]. The first category includes massively parallel processing or MPP Data warehouses that are designed to store huge amount of structured data across a cluster of servers and perform parallel computations over it. Most of these solutions follow shared nothing architecture which means that every node will have a dedicated disk, memory and processor. All the nodes are connected via high speed networks. As they are designed to hold structured data so there is a need to extract the structure from the data using an ETL tool and populate these data sources with the structured data.

These MPP Data Warehouses include:

1) MPP Databases: these are generally the distributed systems designed to run on a cluster of commodity servers. E.g. Aster nCluster, Greenplum, DATAllegro, IBM DB2, Kognitio WX2, Teradata.

2) Appliances: a purpose-built machine with preconfigured MPP hardware and software designed for analytical processing. E.g. Oracle Optimized Warehouse, Teradata machines, Netezza Performance Server and Sun's Data Warehousing Appliance.

3) Columnar Databases: they store data in columns instead of rows, allowing greater compression and faster query performance. E.g. Sybase IQ, Vertica,

InfoBright Data Warehouse, ParAccel. Most of them provide SQLs and UDFs to process the data.

Another category includes distributed file systems like Hadoop to store huge unstructured data and perform Map Reduce computations on it over a cluster built of commodity hardware. The purpose of this paper is to propose big data platform for large-scale data analysis by using Map Reduce framework on unstructured data stored in Gluster file system over scale-out NAS. The rest of the paper is organized as follows: In section 2, we present related work and explain background theory such as Big Data and Big Data Analytics in section 3. In section 4, we introduce our proposed big data platform and then conclusion is described in section 5.

2. Related Work

We survey some of existing big data platforms for large scale data analysis. There are many types of vendor products to consider for big data analytics. More recently, vendors have brought out analytic platforms based on MapReduce, distributed file system, and no-SQL indexing. IBM offers a platform for big data including IBM InfoSphere Biginsights and IBM InfoSphere Streams. IBM InfoSphere Biginsights represents a fast, robust, and easy-to-use platform for analytics on Big Data at rest. IBM InfoSphere Streams is a powerful analytic computing platform that delivers a platform for analyzing data in real time with micro-latency [2]. The Vertica Analytics Platform offers a robust and ever growing set of Advanced In-Database Analytics functionality. It has a high-speed, relational SQL database management system (DBMS) purpose-built for analytics and business intelligence. It offers a shared-nothing, Massive Parallel Processing (MPP) column-oriented architecture [4].

ParAccel Analytic Database (PADB), the world's fastest, most cost-effective platform for empowering analytics-driven businesses. When combined with the WebFOCUS BI platform, ParAccel enables organizations to tackle the most complex analytic challenges and glean ultra-fast, deep insights from vast volumes of data. Netezza, a leading developer of combined server, storage, and database appliances designed to support the analysis of terabytes of data and provide companies with a powerful analytics foundation that delivers maximum speed, reliability, and scalability [5]. 1010data offers a data and analytics platform that is the only complete approach to performing the deepest analysis and getting the maximum insight directly from raw data, at a fraction of the cost and time of any other solution [6]. EMC Greenplum is driving the future of data warehousing and analytics with breakthrough products including the Greenplum Data Computing Appliance, Greenplum Database, Greenplum HD enterprise-ready Apache Hadoop, and Greenplum Chorus.

The SAND Analytic Platform is a columnar analytic database platform that achieves linear data scalability through massively parallel processing (MPP), breaking the constraints of shared-nothing architectures with fully distributed processing and dynamic allocation of resources [9]. Pavlo et al. [8] described and compared MapReduce paradigm and parallel DBMSs for large scale data analysis and defined a benchmark consisting of a collection of tasks to be run on an open source version of MR as well as on two parallel DBMSs.

3. Background Theory

This section provides an overview of big data, big data Analytics, big data storage and big data solution. Due to space constraints, some aspects are explained in a highly simplified manner. A detailed description of them can be found in [1] [2] [9].

3.1 Big Data

The term Big Data applies to information that can't be processed or analyzed using traditional processes or tools. Increasingly, organizations today are facing more and more Big Data challenges. They have access to a wealth of information, but they don't know how to get value out of it because it is sitting in its most raw form or in a semistructured or unstructured format; and as a result, they don't even know whether it's worth keeping. There are three characteristics of Big Data: volume, variety, and velocity.

1) Volume**:** Volume is the first and most notorious feature. It refers to the amount of data to be handled. Many organizations are producing large amounts of data internally, or gathering other large amounts of data from the exterior.

2) Variety: The variety of data that organizations collect has increased in several ways: there are more internal systems having (primarily structured) data collected from them, and there is the rise of internal and external sources of data from semi- or nonstructured social media sources, such as blogs and tweets, as well as data coming from sensors and even plain-text documents.

3) Velocity: As with traditional types of solutions (e.g., the data warehouse), latency periods are being reduced. Information is often sensitive and needs to be used and moved according to certain timeframes to obtain the best possible value from it. Real-time or near real-time answers are common needs in modern organizations.

There are two types of big data: data at rest (e.g. collection of what has streamed, web logs, emails, social media, unstructured documents and structured data from disparate system) and data in motion (e.g. twitter/facebook comments, stock market data and sensor data).

3.2 Big Data Analytics

Big Data analytics is an area of rapidly growing diversity. It emerges from the following reasons:

1) The perception that traditional data warehousing processes are too slow and limited in scalability,
2) The ability to converge data from multiple data sources, both structured and unstructured and
3) The realization that time to information is critical to extract value from data sources that include mobile devices, RFID, the web and a growing list of automated sensory technologies.

Big data analytics is the application of advanced analytic techniques to very big data sets. Advanced analytics is a collection of techniques and tool types, including predictive analytics, data mining, statistical analysis, complex SQL, data visualization, artificial intelligence, natural language processing, and database methods that support analytics (such as MapReduce, in-database analytics, in-memory database, columnar data stores).

Big Data analytics requires massive performance and scalability- common problems that old platforms can't scale to big data volumes, load data too slowly, respond to queries too slowly, lack processing capacity for analytics and can't handle concurrent mixed workloads. There are two main techniques for analyzing big data: the store and analyze approach, and the analyze and store approach [10].

3.2.1 Store and Analyze Approach

The store and analyze approach integrates source data into a consolidated data store before it is analyzed. This approach is used by a traditional data warehousing system to create *data analytics*. In a data warehousing system, the consolidated data store is usually an enterprise data warehouse or data mart managed by a relational or multidimensional DBMS. The advantages of this approach are improved data integration and data quality management, plus the ability to maintain historical information. The disadvantages are additional data storage requirements and the latency introduced by the data integration task.

Two important big data trends for supporting the store and analyze approach are relational DBMS products optimized for analytical workloads (often called analytic RDBMSs, or ADBMSs) and non-relational systems (sometimes called NoSQL systems) for processing unstructured data. A non-relational system can be used to produce analytics from big data, or to preprocess big data before it is consolidated into a data warehouse. Certain vendors in the search and content management marketplaces also use the store and analyze approach to create analytics from index and content data stores.

3.2.2 Analyze and Store Approach

The **analyze and store** approach analyzes data as it flows through business processes, across networks, and between systems. The analytical results can then be published to interactive dashboards and/or published into a data store (such as a data warehouse) for user access, historical reporting and additional analysis. This approach can also be used to filter and aggregate big data before it is brought into a data warehouse. There are two main ways of implementing the analyze and store approach:

1) **Embedding the analytical processing in business processes.** This technique works well when implementing business process management and service oriented technologies because the analytical processing can be called as a service from the process workflow. IBM supports this style of processing in its WebSphere product set. This technique is particularly useful for monitoring and analyzing business processes and activities in close to real-time – action times of a few seconds or minutes are possible here. The *process analytics* created can also be published to an operational dashboard or stored in a data warehouse for subsequent use.
2) **Analyzing streaming data** as it flows across networks and between systems. This technique is used to analyze data from a variety of different (possibly unrelated) data sources where the volumes are too high for the store and analyze approach, sub-second action times are required, and/or where there is a need to analyze the data streams for patterns and relationships. To date, many vendors have focused on analyzing event streams (from trading systems, for example) using the services of a complex event processing (CEP) engine, but this style of processing is evolving to support a wider variety of streaming technologies and data. IBM's InfoSphere Streams product, for example, supports a stream processing engine that creates *stream analytics* from many types of streaming data such as event, video and GPS data. The benefits of the analyze and store approach are fast action times and lower data storage overheads because the raw data does not have to be gathered and consolidated before it can be analyzed.

3.3 Big Data Storage

3.3.1 Big Data Storage for Big Data Analytics

There is a case to be made for shared storage in Big Data analytics. Yet storage vendors and the storage community in general, have yet to make that case to practitioners of Big Data analytics. An example can be seen in the integration of the ParAccel’s Analytic Database (PADB) with NetApp SAN storage. Developers of data storage technology are moving away from expressing storage as a physical device and toward the implementation of storage as a more virtual and abstract entity. As a result, the shared storage environment can and should be seen by Big Data practitioners as one in which they can find potentially valuable data services, such as:

1) **Data protection and system availability:** Storage-based copy functions that don’t require database quiescence can create restartable copies of data to recover from system failures and data corruption occurrences.
2) **Reduced time to deployment for new applications and automated processes:** Business agility is enhanced when new applications can be brought online quickly by building them around reusable data copies.
3) **Change management:** Shared storage can potentially lessen the impact of required changes and upgrades to the online production environment by helping to preserve an “always-on” capability.
4) **Lifecycle management:** The evolution of systems becomes more manageable and obsolete applications become easier to discard when shared storage can serve as the database of record.
5) **Cost savings:** Using shared storage as an adjunct to DAS in a shared-nothing architecture reduces the cost and complexity of processor nodes.

Each of the above mentioned benefits can be mapped to shared-nothing analytics architectures [14].

3.3.2 Big Data Storage Considerations

Many organizations are struggling to deal with increasing data volumes, and big data simply makes the problem worse. To solve this problem, organizations need to reduce the amount of data being stored and exploit new storage technologies that improve performance and storage utilization. From a big data perspective there are three important directions:

1) **Reducing data storage requirements** using data compression and new physical storage structures such as columnar storage.
2) **Improving input/output (I/O) performance** using solid-state drives (SSDs).
3) **Increasing storage utilization** by using tiered storage to store data on different types of devices based on usage [10].

3.4 Big Data Solution

A big data solution provides the technical means to perform operations with high volumes of data in a short period of time, with the ability to treat various types of data from disparate sources. Examples of big data solutions are multi-channel customer sentiment and experience analysis, detect life-threatening conditions at hospitals in time to intervene, make risk decisions based on real-time transactional data, identify criminals and threats from disparate video, audio, and data feeds and predict weather patterns to plan optimal wind turbine usage, and optimize capital expenditure on asset placement.

The big data solution can provide these abilities. They are:

1) **Deep Analytics** — a fully parallel, extensive and extensible toolbox full of advanced and novel statistical and data mining capabilities
2) **High Agility** — the ability to create temporary analytics environments in an end-user driven, yet secure and scalable environment to deliver new and novel insights to the operational business
3) **Massive Scalability** — the ability to scale analytics and sandboxes to previously unknown scales while leveraging previously untapped data potential
4) **Low Latency** — the ability to instantly act based on these advanced analytics in your operational, production environments [12].

4. Proposed System Architecture

4.1 Big Data Storage Architecture

Traditional models are used to dealing with text and numbers which lend themselves to traditional database models and processing techniques. Today’s big data include multiple object types. These requirements have broken the traditional data storage models and created the need for new architectures to effectively store and deliver this data to the analytics systems that do the analysis. Older storage architectures couldn’t scale to the size required or hold the diverse data types that are being created. Limitations on the amount of data that could be stored in an array was in the 100s of terabytes range but the file systems they provided could not scale beyond 16 terabytes [7].

These older architectures used the fundamental approach of scale up vs. scale out. The primary difference is how the system uses resources. Scale up system would provide a small number of access points, or data servers, that would sit in front of a set of disks protected with RAID. As these systems needed to provide more data to more users the storage administrator would add more disks to the back end but this only caused to create the data servers as a choke point. Larger and faster data servers could be created using faster processor and more memory but this architecture still had significant scalability issues.

Scale out uses the approach of more of everything—instead of adding drives behind a pair of servers, it adds servers each with processor, memory, network interfaces and storage capacity. This architecture required a number of things to make it work from both a technology and financial aspect. Some of these factors include Clustered architecture, Distributed/parallel file system and Commodity hardware. There are a number of significant benefits to these new scale-out systems that meet the needs of big data challenges. They are manageability, elimination of stovepipes, just in time scalability and increased utilization rates. Figure 1 describes the large scale shared storage system architecture for big data.

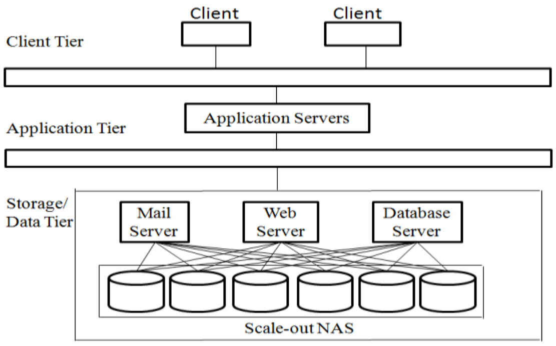

Figure 1. Large scale shared storage system architecture for big data

4.2 Proposed Big Data Platform

In general for big data analytics, there are three approaches: 1) direct analytics over massively parallel processing data warehouses, 2) indirect analytics over hadoop and 3) direct analytics over hadoop. The proposed approach performs analytics over Hadoop MapReduce framework and Gluster file system. All the queries for analytics are executed as Map Reduce jobs over big unstructured data placed into Gluster file system. By using this approach, a highly scalable, fault tolerant and low cost big data solution can be achieved. Figure 2 describes the proposed big data approach.

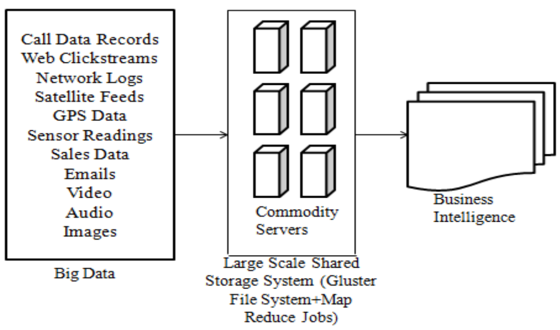

Figure 2. Proposed big data approach

By combining scalability to petabytes and beyond, affordability (use of commodity hardware), flexibility (deploy in any environment), linearly scalable performance, high availability, unified files and objects, file system for apache hadoop and superior storage economics of Gluster file system with parallel data processing, schema free processing and simplicity of Map Reduce programming model, its open source implementation Hadoop and other Hadoop open source projects (pig, hive, jaql), the proposed big data platform can perform large scale data analysis efficiently and effectively. Proposed big data platform is shown in Figure 3.

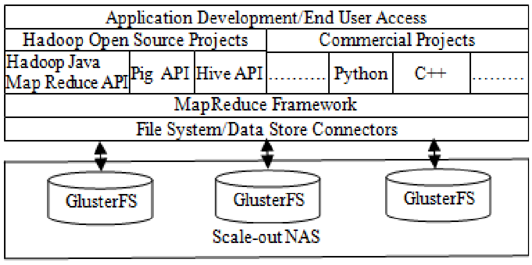

Figure 3. Proposed big data platform

4.2.1 Hadoop and MapReduce Framework

Data growth – particularly of unstructured data – poses a special challenge as the volume and diversity of data types outstrip the capabilities of older technologies such as relational databases. Organizations are investigating next generation technologies for data analytics. One of the most promising technologies is the Apache Hadoop software and MapReduce framework for dealing with this “big data” problem.

A MapReduce framework typically divides the input data-set into independent tasks which are processed by the map tasks in a completely parallel manner. The framework sorts the outputs of the maps, which are then input to the reduce tasks. Typically both the input and the output of the jobs are stored in a file-system. The framework takes care of scheduling tasks, monitoring them and reexecuting the failed tasks [13].

4.2.2 Gluster File System

Gluster file system is a scalable open source clustered file system that offers a global namespace, distributed front end, and scales to hundreds of petabytes without difficulty. It is also a software-only, highly available, scalable, centrally managed storage pool for unstructured data. It is also scale-out file storage software for NAS, object, big data. By leveraging commodity hardware, Gluster also offers extraordinary cost advantages benefits that are unmatched in the industry. There are many advantages of Gluster over any other file systems. These advantages are:

1) It is faster for each individual operation because it calculates metadata using algorithms and that

approach is faster than retrieving metadata from any storage media.
2) It is faster for large and growing individual systems because there is never any contention for any single instance of metadata stored at only one location.
3) It is faster and achieves true linear scaling for distributed deployments because each node is independent in its algorithmic handling of its own metadata, eliminating the need to synchronize metadata.
4) It is safer in distributed deployments because it eliminates all scenarios of risk which are derived from out-of-synch metadata.

Gluster file system can be used in place of HDFS, which brings all that software-based data protection and functionality to the Hadoop cluster and removes the single point of failure issue. GlusterFS 3.3 beta 2 includes compatibility for Apache Hadoop and it uses the standard file system APIs available in Hadoop to provide a new storage option for Hadoop deployments. Existing MapReduce based applications can use GlusterFS seamlessly. This new functionality opens up data within Hadoop deployments to any file-based or object-based application [3]. Figure 4 describes Gluster file system compatibility for Apache Hadoop.

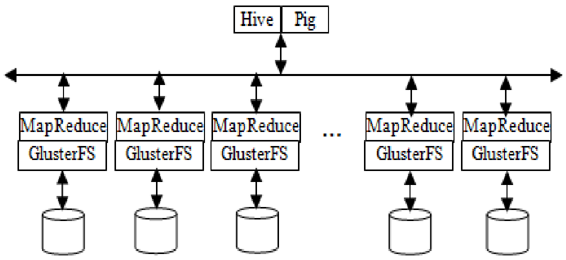

Figure 4. GlusterFS compatibility for Apache Hadoop

The following are the advantages of Hadoop Compatible Storage with GlusterFS: 1) Provides simultaneous file-based and object-based access within Hadoop, 2) Eliminates the centralized metadata server, 3) Provides compatibility with MapReduce applications and rewrite is not required and 4) Provides a fault tolerant file system. Hadoop integration with Gluster is shown in Figure 5.

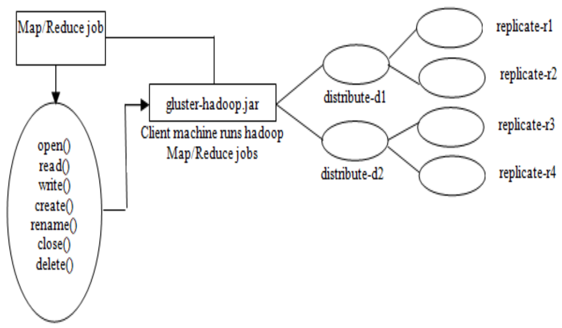

Figure 5. Hadoop integration with Gluster

5. Conclusion

Big Data is a growing problem for corporations as a result of sheer data volume along with radical changes in the types of data being stored and analyzed, and its characteristics. The main challenges of big data are data variety, velocity and volume, and analytical workload complexity and agility. To address these challenges, many vendors have developed big data platforms. In this paper, we have proposed big data platform for large scale data analysis by using Hadoop Map Reduce Framework and Gluster file system over scale out NAS. But Hadoop/MapReduce is batch-like, and not immediately suitable for real time analysis, unsuited to ad hoc queries. Hadoop solves volume and variety issues and so we still need to solve velocity issue.

References

[1] P. Carter, “Big Data Analytics: Future Architectures, Skills and Roadmaps for the CIO”, IDC, September 2011.
[2] C. Eaton, T. Deutsch, D. Deroos, G. Lapis and P. Zikopoulos, “Understanding Big Data: Analytics for Enterprise Class Hadoop and Streaming Data”,McGraw-Hill ,2011.
[3] Gluster, “Gluster File System 3.3-Beta 2 Hadoop Compatible Storage”, August 2011.
[4] Hewlett-Packard, “HP Advanced Information Services for the Vertica Analytics Platform”, January 2012.
[5] IBM Software, “IBM Netezza Analytics”, Information Management, August 2011.
[6] Information Builders, “Advanced Business Analytics and the New Era of Data Warehousing”, March 2011.
[7] M. Janson, “Big Data and the New Storage Architecture”, OnX Enterprise Solutions, November 2011.
[8] A. Pavlo, E. Paulson and A. Rasin,” A Comparison of Approaches to Large-Scale Data Analysis,” in Proceedings of the 35th SIGMOD International Conference on Management of Data. ACM, 2009.
[9] P. Russom,”Big Data Analytics”, TDWI best practices report, fourth quarter 2011.
[10] C. White, “Using Big Data for Smarter Decision Making”, BI Research, July 2011.
[11] http://bigdataanalytics.blogspot.com
[12] http://blogs.oracle.com/datawarehousing/entry/big_data_achieve_the_impossible
[13] http://hadoop.apache.org/mapreduce
[14] http://searchstorage.techtarget.com/feature/Understanding-Big-Data-analytics