

A Novel set of tools to “Extract, Rework, Batch” [ERB] Big Data with Apache Hadoop Environment

Ketavarapu Nageswara Rao¹, Maddineni Karuna Sree², K.Narasimha Sastry³,
Duvvada Rajeswara Rao⁴

¹(CSE M tech /KL University, India) ²(CSE M tech/KL University, India)

³(CSE M tech/KL University, India) ⁴(CSE Prof/KL University, India)

Abstract: Over the previous couple of years, organizations across public and personal sectors have created a strategic call to show massive information into competitive advantage. The challenge of extracting price from huge information is comparable in many ways to the antique drawback of distilling business intelligence from transactional knowledge. At the center of this challenge is that the method costumed extract information from multiple sources, remodel it to suit your analytical wants, and load it into an information warehouse for consequent analysis, a process called “Extract, Rework, Batch” (ERB). The character of huge knowledge needs that the infrastructure for this method can scale cost-effectively. Apache Hadoop has emerged because the factual normal for managing huge data. This paper examines some of the platform hardware and software considerations in using Hadoop for (ERB).

Key words: Bigdata, Hadoop, Map Reducer, ERB, HDFS.

I. Introduction

ERB holdup in Big Data Analytics

Big Data deals with the huge amounts, at least terabytes, of poly-structured data that flows incessantly through and around firms, as well as sensor logs, texts, videos, and transaction able records. The trade benefits of analyzing this data can be considerable. Current study by the MIT Sloan School of Management ^[1], organizations that apply analytics are twice as feasible to be top performers in their trade as those that don't. commerce analysts at a large corporation such as Intel, for example, with its worldwide market and difficult supply chain, have long required imminent into customer demand by analyzing remote data points culled from market information and trade transactions. progressively the records we need is entrenched in pecuniary reports, deliberation forums, social networks, weather condition reports, news reports, wikis, tweets, and pages, as well as transactions. By resolving all the data presented, decision-makers can enhance assess competitive coercion, foresee changes in user behavior, make stronger supply chains, improve the effectiveness of advertising campaigns, and increase business stability.

Lots of of these benefits are not latest to organizations that have established processes for incorporating business intelligence (BI) ^[2] and analytics into their executive. though, most organizations have yet to take full lead of new technologies for using big data. set simply, the cost of the technologies required to store and evaluate large volumes of miscellaneous data has dropped, gratitude to open source software operating on industry-standard hardware. The cost has droned so much, in fact, that the key tactical question is no longer what data is significant, but relatively how to extract the most value from all the data available. firstly swilling, storing, and processing big data requires a cost-effective infrastructure that can level with the amount of statistics and the extent of analysis. The majority of the organizations with conventional data platforms classically relational database management systems (RDBMS) united to enterprise data warehouses (EDW) using ERB tackles uncover that their inheritance infrastructure is either precisely incapable or financially unfeasible for storing and analyzing big data. A conventional ERB process extracts data from various sources, then cleanses, formats, and batches it into a data warehouse for scrutinizing. while the source data sets are huge, speedy, and unstructured, conventional ERB can be converted into the bottleneck, as it is too compound to develop, too expensive to operate, and takes too extensive to execute. nearly 80 % of the development effort in a big data project goes into data addition and only 20 % goes to data analysis. in addition, a conventional (EDW) platform may cost upwards of USD 60K per terabyte. Analyzing one peta byte the quantity of data Google exercises in 1 hour would cost USD 60M. noticeably “much of the same” is not a big data tactic that any CIO can afford. So, pierce into Apache Hadoop.




Fig1: ERB Process

II. Apache Hadoop For Big Data

When Yahoo, Google, Facebook, and other companies wide-ranging their services to web scale, the amount of data they collected usually from user communications online would have besieged the capabilities of traditional IT framework. So they built their own. In the concern of advancing the expansion of foundation infrastructure components rapidly, Apache Hadoop has rapidly emerged as the de facto customary for running large volumes of unstructured records

Apache Hadoop is an open source disseminated software platform for storing and dispensing data. Written in Java, it runs on a group of industry customary servers configured with direct-attached storage. by Hadoop, you can pile up petabytes of data dependably on tens of thousands of servers while scaling performance economically by simply adding inexpensive nodes to the cluster.

Central to the scalability of Hadoop is the distributed processing framework recognized as MapReduce (Figure 2). MapReduce helps programmers answer data-parallel problems for which the data set can be subdivided into small parts and processed autonomously.

MapReduce is a main advanced because it allows regular developers, not just those skilled in high-performance computing, to use parallel programming constructs exclusive of worrying about the complex details of intra-cluster announcement, task monitoring, and failure handling. MapReduce simplifies all that. The classification splits the input data-set into numerous chunks, each of which is allotted a map task that can develop the data in parallel. every map task reads the input as a set of (key, value) pairs and produces a transformed set of (key, value) pairs as the output. The structure shuffles and sorts outputs of the map tasks, transfer the intermediate (key, value) pairs to the reduce tasks, which batch them into end results. MapReduce makes use of JobTracker and TaskTracker mechanisms to schedule total tasks, monitor them, and start over any that fail. The Apache Hadoop platform also includes the Hadoop Distributed File System (HDFS), which is designed for scalability and fault tolerance. HDFS stores bulky files by dividing them into blocks (usually 64 or 128 MB) and replicating the blocks on three or more servers. HDFS provide APIs for MapReduce applications to read as well as write data in parallel. capability and performance can be scaled by adding Data Nodes, and a single Name Node device manages data placement and monitors server accessibility. HDFS clusters in production use today dependably hold peta bytes of data on thousands of nodes.

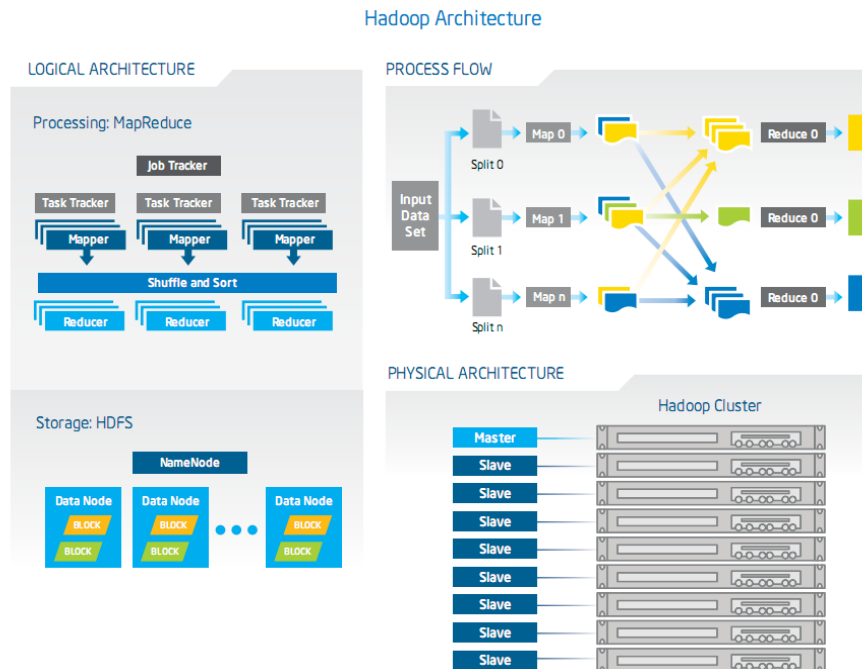


Fig 2: MapReduce, the encoding archetype implemented by Apache Hadoop, splits a batch job into various smaller tasks for parallel processing on a distributed system. HDFS, the distributed file system stores the data consistently.

Other than MapReduce and HDFS, Apache Hadoop has many other components, some among them are very useful for ERB.

1.1 Apache Flume is a strewn system for collecting, aggregating, and also touching massive amounts of data from various sources into HDFS or another crucial data store. Enterprises typically gather log files on application servers or other systems and records the log files in order to obey with regulations. Being capable to ingest and evaluate that unstructured or semi-structured data in Hadoop can turn this imprudent resource into an essential asset.

1.2 Apache Sqoop is a tackle for transferring data among Hadoop and relational databases. You can use Sqoop to bring in data from a MySQL or Oracle database into HDFS, run MapReduce on the data, and after that export the data back to RDBMS. Sqoop automatize these processes, via MapReduce to import and export the records in parallel by fault-tolerance.

1.3 Apache Pig and Apache Hive are programming languages that make simpler the development of applications retaining the MapReduce framework. HiveQL is a language of SQL and supports a rift of the syntax. even though slow, Hive is being actively improved by the developer community to facilitate low-latency queries on HDFS and Apache H Base and. Pig Latin is a procedural programming language ttat facilitates high-level abstractions for MapReduce. can extend it by User Defined Functions written in Java, Python, and some other languages.

1.4 JDBC/ODBC Connectors meant for H Base and Hive are mostly proprietary components integrated in distributions meant for Apache Hadoop software. They offer connectivity with SQL applications with translating customary SQL queries into HiveQL commands that can be executed ahead the data in HDFS or HBase Hadoop is a dominant platform meant for big data storage and organizing. though, its extensibility and innovative renew questions about data integration, data eminence, authority, security, and a congregation of other issues that firm with established BI processes have extensive use for granted. regardless of many summons of incorporating Hadoop into a conventional BI environment, ERB has confirmed to be a frequent use case for Hadoop in firm. therefore what details its popularity?

2. ERB, EBR, AND ERBT WITH APACHE HADOOP

ERB tackle move data from one place to some other place by implementing three functions:

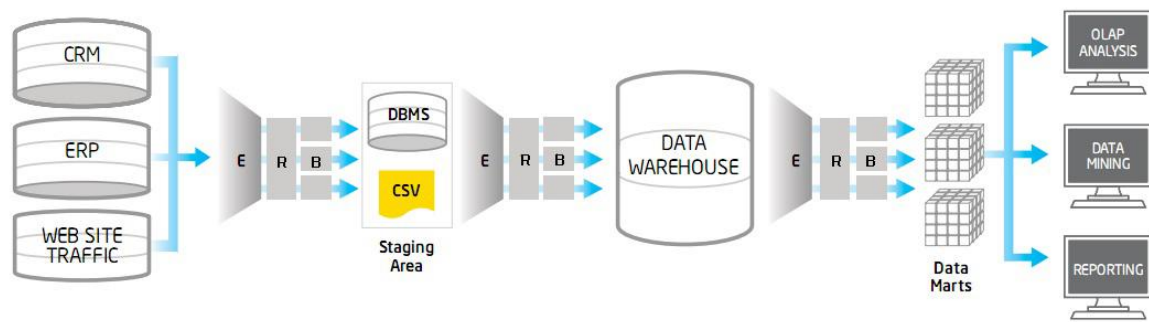


Figure 3. The traditional ERB architecture

2.1 Extract records from sources like ERP or CRM applications.

At the extract step, you might need to collect records from numerous source systems and in various file formats, like flat files with measures (CSV) and XML files. You might also need to gather data from inheritance systems that pile up data in esoteric formats no one else uses any longer. This sounds simple, but can in fact be one among the key obstacles in getting an ERB elucidation off the ground.

2.2 rework that data into a regular format that fits further most data in the warehouse.

The reworking step may embrace multiple data handling, such as moving, dividing, translating, including, organization, resolving, and more. For instance, a customer name may be divided into first and last names, or dates might be altered to the standard ISO format (e.g., from 09-06-14 to 2014-09-08). frequently this step also involves legalizing the data beside data quality rules.

2.3 batch the data into the data warehouse for scrutinizing.

This step can be done in batch processes or row by row, more or less in real time.

In the beginning days, before ERB tackle existed, the only way to incorporate data from diverse sources was to hand-code scripts in languages such as COBOL, PL/SQL, and RPG. obsolete though it seems, nearly 45 % of all ERB work still today is done by such safe coded programs. though they are error-prone, slow to build up, and rigid to maintain, they have reliable users who seem inflexible to the charms of ERB tackle, like Oracle Warehouse Builder,* that generate code for databases.

Code generators also have boundaries, as they work with only a limited set of databases and are frequently grouped with them. In contrast, the further generation of ERB tackles includes a general-purpose engine that implements the alteration and a set of metadata that stores the renovation logic. as engine-based ERB tackle such as Informatica Power center and Pentaho Kettle are autonomous of source and target data stores, they are much adaptable than code generating tools.

A conventional ERB design (Figure 3) accommodates various ERB iterations as well as an intermediary step, performed in the “staging area,” that gets data from the source systems as fastly as possible. A staging area may makes use of database or plain CSV files, which enables the process rapid than inserting data into a database table. Some Additional ERB iterations may be performed to transfer data between EDW and data marts, and this supports particular analytic purposes and end-user tackle. a lot has altered in data warehousing from the past two decades. primarily, databases have become immensely more powerful. Presently RDBMS engines support compound transformations in SQL, together with in-database data mining, in-database data quality, profiling, cleansing, , statistical algorithms, validation, hierarchical and drill-down functionality, and a lot more. It has become more capable to carry out most types of “transformation” in the RDBMS engine. Figure 2. The conventional ERB design has served firms well for the past twenty years and so many deployments use it still.

As a result, EBR came out as an substitute approach in which data is extracted from the sources, batched into the target database, and then altered and incorporated into the desired format. inside the target database All the heavy data processing takes place. The benefits’ of this approach is that a database system is best suited to handle massive workloads for which thousands of millions of records required to be incorporated. RDBMS engines are also suitable for disk Input and output that increases throughput. the system performance scales with the RDBMS hardware.

But ERB is not dead still. Conventional ERB traders have improved their tackles with pushdown SQL abilities in which alteration can take place either in the ERB engine or after batching inside the database. The result was ERBT approach, supported regularly well by primary database traders like Oracle (Oracle Warehouse Builder) and Microsoft (SQL Server Integration Services).

None of these solutions is worthless or easy, and their cost and difficulty are intensified by big data. Regarding eBay, this in 2011 had over 200 million products for sale, divided into 50,000 categories, and are bought and sold by 100 million registered customers of which entailed approximately 9 petabytes of data. Google apparently processes over 24 petabytes of data each day. AT&T processes 19 petabytes by their networks every day, the video game World of War craft makes use of 1.3 petabytes of storage space. All these figures are previously out-of-date of this writing as online data is emerging most frequently Under these situations, Hadoop brings minimum two major utilities to conventional ERBT:

2.4 Ingest enormous amounts of data without mentioning any schema

on write. A key attribute of Hadoop is defined “no schema on write down,” this means you don’t require to declare the data schema before piling data in to Hadoop. This is correct for structured data, social media data, and unstructured data . despite of your incoming data have unambiguous or embedded structure, you can frequently load it into Hadoop as it-is, where it is offered for write downstream analytic processes.

2.5 Offload the renovation of raw data by parallel processing at range.

when the data is in Hadoop ,you can carry out the conventional ERB tasks of cleaning, normalizing, aligning, assembling data for your EDW by retaining the huge extensibility of MapReduce. Hadoop by off-loading the ingestion, transformation, and integration of unstructured data into your data warehouse (Figure 3) allows you to avoid the transformation bottleneck in your conventional ERBT. as Hadoop permits you to cuddle more data types than before, it enhance your data warehouse in ways that will be unimaginable or unaffordable. you can considerably speed up the ERBT jobs Because of its expandable performance. as data stored in Hadoop can continue for a much longer duration, you can offer much more rough, detailed data by your EDW for high-fidelity breakdown. By Using Hadoop in this manner, the firm gains some extra ability to store and utilize data that they may require, data that might never be encumbered into the data warehouse. Consider an example, data analysts might need to use the bulk amounts of source data from web logs, third-party stores, or social media, stored on Hadoop to improve new systematic models that make investigation and innovation. In hadoop the data can be stored efficiently and recover it as needed by making use of Hive or some other logical tools subject to the platform, without disturbing the EDW environment.

in spite of whether your organization takes the ERB, EBR, or ERBT approach to data warehousing, you can deduce the functional cost of your complete BI/DW solution by unloading familiar revolution pipelines to Apache Hadoop and on HDFS making use of tolerant platform for dealing with massive amounts of assorted data.

III. Selecting The Physical Infrastructure For ERB With Hadoop

The principle of thumb for Hadoop infrastructure development has been long to “fling more nodes at the difficulty.” This is a logical advance when the range of your cluster equals a web-scale confront, like returning search outcomes or customizing web pages for thousands of millions of online customers. But a in an organization a classical Hadoop cluster have nearly hundred nodes and is encouraged by Information Technology resources that are considerably much more inhibited than those of Facebook and Yahoo! . companies running these clusters must remain to the capacity schedule and performance regulation processes representative of other Information Technology infrastructures. Sooner, they require stipulate, construct, and refrain their Hadoop cluster for the particular workloads their work run. Workloads differ extensively, so it is essential to select and compute, network and storage, software infrastructure to match with specific requirements. But, let’s scrutinize the common misapprehension that all Hadoop workloads are input and output -bound before taking into account each of these cases., developers at Intel have been testing the performance of successive releases of Hadoop For the past 3 years on consecutive generations of Intel processor-dependend servers utilizing a orientation set of test workloads. So Many of these workloads, in cumulative, look like real-world applications, like analytics and ERB. Considering an in-depth expedient of the test clusters, we founded that Input, output and CPU utilization differ extensively across workloads and within the states of MapReduce in each and every workload.

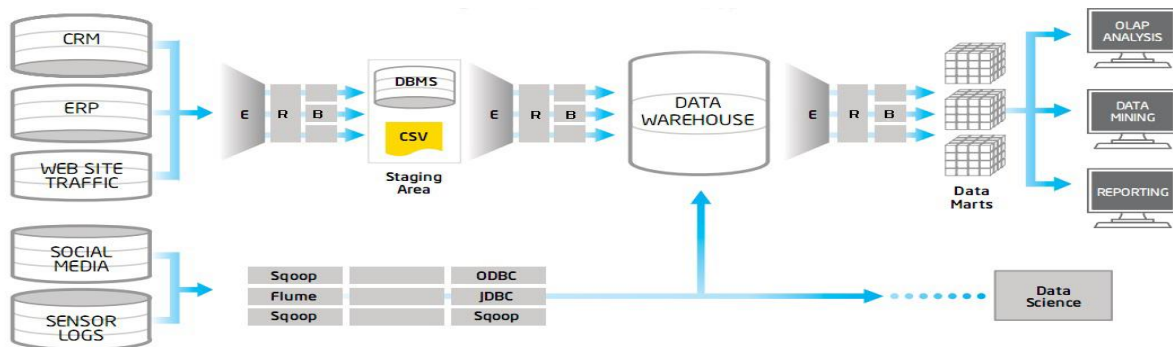


Figure 4. Offload ERB with Hadoop

3.1 **TeraSort** (Reduce: I/O Bound Map: CPU Bound;) TeraSort alters data from one illustration to another. TeraSort tends to be Input and Output bound as the size of data does not alter from input throughout the mix up to output., TeraSort shows very massive CPU utilization and restrain able disk Input and output during the shuffle and map phases when we compact map output so as to lessen the disk and network Input and output at the time of shuffle phases, and modest CPU utilization and powerful disk Input and Output at the reduce phase.

3.2 **Word Count** (CPU Bound) Word Count obtains a small amount of remarkable data from a massive data set, which means that the reduce output and map output are very much slighter than the job input. As an outcome, the CPU vault with high CPU utilization and light disk and network I/O is called as Word Count .

3.3 **Notch Indexing** (Reduce: I/O Bound Map: CPU Bound); The input to this workload is nearly 2.4 million web pages produced by crawling Wikipedia. Notch Indexing ease the crawl data in the map level, which is nothing but CPU bound, and transforms transitional results to reversed index files in the reduce level, which is called I/O bound.

3.4 **PageRank Search** (CPU Bound) The Page Rank workload is delegate of the original motivation for Google search depended on the Page Ranking algorithm for link analysis. This workload for iterations of several jobs spends most of th time, and these jobs are CPU bound, with low to medium memory utilization and disk I/O.

3.5 **Bayesian Classification** (I/O Bound) This workload equips the instructor part of the Naïve Bayesian classifier, which is a popular algorithm for discovering knowledge and data mining. The workload contains 4 chained Hadoop jobs, which are typically disk I/O bound, excluding the map tasks of the 1st job, which has the high CPU utilization and this consumes most of the time.

3.6 **K-means Clustering** (I/O Bound in clustering CPU Bound in iteration;) This workload 1st solves the central part of each cluster by organizing a Hadoop job monotony till iterations congregate or the maximum no.of iterations is reached. This is called CPU-bound. neatly it runs a clustering job which assigns each and every sample to a cluster.

Even a rapid scan of different delegate workloads shows that the system resource consumption of Hadoop is more significant for firms than it has been for first adopters and more tough than many have understood. Intel has worked intimately with a number of consumers to help build up and organize a balanced platform for different real-world organizes of Hadoop. Our evaluations and standard efforts have led us to make numerous recommendations when consumers regard infrastructure hardware and dissemination of Apache Hadoop software.

IV. Compute

The Intel Xeon processor E5 family facilitates a strong platform for many Hadoop workloads. A number of characteristics constructed into the Intel Xeon processor E5 family are predominantly well matched for Hadoop. One characteristic is Intel Integrated Input and Output, which decreases Input and Output latency by up to 32 % and increases Input and Output bandwidth by nearly 2x.2,3 Another one is Intel Data Direct Input and Output Technology (DDIO), which facilitates Intel® Ethernet adapters to converse directly with processor cache, than only with main memory. This characteristic delivers more I/O bandwidth and lesser latency, which is mainly favorable when dealing with large data sets. Another characteristic that has seized the attention of Hadoop end user is Intel Advanced Encryption Standard New Instructions (AES-NI), which speeds up general cryptographic functions to obliterate the presentation consequence characteristically related with decryption and encryption of big-data files. In so many surveys organized by Gartner, more than seventhly percent of CIO’s outlined that cooling and power issues were the greatest confront they encountered in the data center.4 improving energy-efficiency can be predominantly important for Hadoop clusters, which grows with data amount. the Intel Atom™ processor With its system on a chip design, power casing as low as 6 watts, provides better energy efficiency and density for some workloads. as Intel Atom processors are found on the industry-standard eighty six instruction set, they are attuned code that runs only Intel Xeon processors. This consistency helps to diminish total cost of ownership by decreasing software development, system management costs and

porting. In addition, Intel provides a thermal and power management software artifact known as Intel® Data Center Manager (Intel DCM). Intel DCM makes use of the instrumentation in Intel Xeon and forthcoming Intel Atom processors and incorporates with offered management helps using standard APIs. can use it to supervise thermal and power data in real time for blades and entity servers, and also for rows and racks, and logical server category. You can restrict power with set confines or workload-correspondent policies, organize alerts for thermal, power and events, and store and study data to recover capacity planning.

4.1 Memory

adequate system memory is necessary for high output of many numbers of parallel MapReduce jobs. Hadoop classically requires 48 Gigabytes to 96 Gigabytes of RAM per server, and 64 Gigabytes is best in most cases. transversely available memory channels constantly balance server memory to shun memory-bandwidth bottlenecks. one of the most common causes of server failure and data corruption are Memory errors , so error-correcting code (ECC) memory is extremely suggested.⁵ ECC is reinforced in all servers regarding Intel Xeon processors and in micro-servers depended on the Intel Atom processor family.

4.2 Storage

Each and every server in a Hadoop cluster requires a comparatively large number of storage drives to shun Input and Output bottlenecks. 2 hard drives per processor core classically deliver good outcomes. though, a single solid-state drive (SSD) per core can deliver superior Input and output throughput, concentrated latency, and enhanced overall cluster concert. Intel® SSD 710 Series SATA SSDs offers considerably faster read and write presentation than mechanical hard drives, this additional performance can be expensive for latency perceptive MapReduce applications. It can also hastens jobs as intermediate consequence files are altered between the reduce and map the phase. Intel tests show that substituting mechanical drives with Intel SSDs can enhance concert by as much as 80 %.⁶ can also use SSDs and mechanical drives together, using Intel® Cache Acceleration Software. This tiered storage copy offers some of the performance advantages of SSDs at a inferior acquisition cost. If you use hard drives, 7,200 RPM SATA drives offer a good balance between performance and cost. Run the drives in the Native Command Queuing (NCQ) with Advanced Host Controller Interface (AHCI) mode enabled to develop presentation when multiple read and write desires are invoked concurrently. though you can use RAID 0 to logically merge smaller drives into a bigger pool, Hadoop repeatedly arranges data attachments and idleness across nodes, so RAID is not suggested.

4.3 Network

A rapid network not only allows data to be exported and imported quickly, but can also perk up presentation for the shuffle phase of MapReduce applications. A 10 Gigabit Ethernet network advances a easy, cost-effective elucidation. Intel tests have shown that making use of 10 Gigabit Ethernet slightly than 1 Gigabit Ethernet in a Hadoop cluster can get better concert for key operations by up to 4x when using predictable hard drives. The presentation enhancement is even better when using SSDs—up to 6x.⁷ The greater enhancement with SSDs can be credited to faster writes into the storage subsystem. As a Hadoop cluster grows to embrace multiple server racks, you can measure network presentation by connecting each and every thing of the ten Gigabit Ethernet rack-level switches to a 40 Gigabit Ethernet cluster-level switch. As requirements continue to increase, you can interconnect so many cluster-level switches and add an uplink to switching infrastructure’s highest level.

4.4 Software

Apache Hadoop is open source software that is freely offered from the apache.org source code depository. A number of test deployments and early adopters download the source directly and build their platform on Apache Hadoop circulation. Firms and other organizations that requires a vendor-supported platform, though, look to one of many independent software vendors (ISVs) to facilitate a complete product with updates, software, and services. The Intel® Distribution for Apache Hadoop platform is an enterprise-grade software platform that contains Apache Hadoop along with some other software mechanism. The Intel Distribution has many numbers of distinctive features.

4.5 Built from the silicon up for security and performance.

Running the Intel Distribution processors enables Hadoop to fully use the security features and performance available in the x86 instruction set in common and the Intel Xeon processor .consider example, the Intel Distribution includes improvements that take pro’s of Intel AES-NI,present in Intel Xeon processors, to speed up cryptographic functions, removing the distinctive performance penalty of decryption and encryption of files in HDFS.

4.6 Automated presentation tuning.

This incorporated mechanism saves time while distributing many configurations. The Intel Distribution has an intellectual mechanism used for configuring the Hadoop cluster. The Intel Active Tuner makes use of a genetic algorithm to try so many parameters and congregate frequently on the optimal configuration for a presented Hadoop application.

4.7 sustain for a extensive range of analytic applications.

These applications comes from privileged in the field, most of them are already collaborating at the hardware layer with Intel. For example, Intel is functioning closely with Dell Katanga, Revolution Analytics, , SAS, SAP and several analytics ISVs to optimize the presentation and measurability of the overall solution.

V. Conclusion

The most recent wave of big data is prompting innovative opportunities and new summons for enterprise across every commerce. The confront of data integration include data from social media and other unstructured data into a conventional BI environment—is one of the most vital issues facing CIOs and IT managers. Apache Hadoop offers a massively scalable and cost-effective and platform for consuming big data and preparing it for scrutiny. Using Hadoop to devolve the conventional ERB processes can decrease time to scrutiny by hours or even days. Running the Hadoop cluster proficiently means selecting an finest infrastructure of servers, networking, storage, and software. Intel provides software as hardware platform mechanism to help you deploy and design an proficient, high-performing Hadoop cluster perfect for big data ERB. Take pros of Intel orientation architectures, professional services, training, and technical support to fasten up your operation and reduce risk.

References:

- [1]. MIT Sloan study; results published in MIT Sloan Management Review. “Big Data, Analytics and the Path From Insights to Value,” Steve LaValle, Eric Lesser, Rebecca Shockley, Michael S. Hopkins, Nina Kruschwitz, December 21, 2010.
- [2]. The claim of up to 32% reduction in I/O latency is based on Intel internal measurements of the average time for an I/O device read to local system memory under idle conditions for the Intel ® Xeon® processor E5-2600 product family versus the Intel® Xeon® processor 5600 series
- [3]. Gartner 2013 CIO Survey. <http://www.gartner.com/technology/cio/cioagenda.jsp>
- [4]. <http://www.intel.com/content/dam/www/public/us/en/documents/white-papers/big-data-apache-hadoop-technologies-for-results-whitepaper.pdf>
- [5]. <http://hadoop.apache.org/docs/current/api/org/apache/hadoop/examples/terasort/package-summary.html>
- [6]. <https://issues.apache.org/jira/browse/HADOOP-692>.
- [7]. www.datastax.com/wp-content/uploads/2012/09/WP-DataStax-HDFSvsCFS.pdf.
- [8]. hpcsystems.com/Why-HPCC/HPCC-vs-Hadoop.
- [9]. <http://en.wikipedia.org/wiki/Hadoop>.
- [10]. techcrunch.com/tag/big-data/ [<http://techcrunch.com/tag/big-data/>]
- [11]. www.greenplum.com/blog/tag/pivotal-hd [<http://www.greenplum.com/blog/tag/pivotal-hd>]



FIRST A. Author received B.Tech degree in information science and technology from Koneru Lakshmaiah College Of Engineering, Vaddeswaram,Guntur,AP in 2011 and pursuing M.Tech degree in Computer Science And Engineering at Koneru Lakshmaiah University.



SECOND B. Author Received B.Tech degree in Computer Science and Engineering from Sarojini Institute Of Technology,Telaprolu,Andhra Pradesh in 2013 and pursuing M.tech degree in Computer Science and Engineering at Koneru lakshmaiah University.



THIRD C. Author received MCA degree from KL University Vaddeswaram Guntur, A.P in 2010 and pursuing M.Tech degree in Computer Science And Engineering at Koneru Lakshmaiah University.



FOURTH D. Author Professor PhD. He had published research papers at National and International Journals, National and International Conferences. At present he is working has Professor at KL University Vijayawada.