

Impact of Big Data: Networking Considerations and Case Study

Yong-Hee Jeon[†]

Catholic University of Daegu, Gyeongsan, Rep. of Korea

Summary

Due to the explosive growth of data volume by mobile devices and SNS(Social Networking Service), Big Data has recently become one of the important issues in the networking world. Big traffic is generated as Big Data processing steps and multiple regionally distributed data centers are included, and/or data are delivered among clusters for the purpose of storage hierarchy management. Therefore, Hadoop clusters in such a Big Data environment require high-speed networking fabric with multi-Giga bits speed. In this paper, networking infrastructure considerations to support Big Data are studied and Big Data networking architecture is presented through the case study of Cisco.

Key words:

Big Data, Big Traffic, Networking Considerations, Case Study.

1. Introduction

Big Data is also called as “Very large data, Extreme data, and Total data, etc.” and the first criterion was the volume of data. Although there is no exact definition of Big Data, it sometimes refers to more data than ZB (Zeta Byte) range and also means data which require distributed parallel processing technology for the analysis of large volume of data such as Hadoop. 1 ZB is a huge amount of data which corresponds to one trillion Giga bytes[1].

As data volume grows explosively by mobile devices and SNS(Social Networking Service), Big Data has recently become an important issue in IT(Information Technology) field. Among them, the scale of data generation by widespread usage of mobile devices is becoming huge. We have already entered in ‘the age of zeta bytes’ as the digital information amount generated by whole world data reached 1.8 ZB in 2011[1].

According to Cisco, it was forecasted that mobile data grows with an average rate of 78% from 2011 to 2016. The only mobile traffic scale generated in 2016 was forecasted to reach 10.8 Exa Bytes[2]. One Exa byte equals to one quintillion bytes (1 ZB = 1,024 EB).

Based on the definitions described above, Big Data may be further defined as a huge amount of structured or unstructured data set that is difficult to collect, store, analyze, and manage with the existing methods due to its volume. By Mckinsey, Big Data is defined as ‘such data

which exceeds the range possible to store, manage, and analyze it by general database software’. On the other hand, it is defined by Korean President’s Council on National ICT Strategies as “an Information Technology which may extract valuable information by utilizing and analyzing a large volume of data, and may actively react and predict change based on the knowledge generated”[2-3]. Therefore, the meaning of Big Data is becoming expanded as relative one which may obtain a value beyond some criteria rather than simple data volumes and technology aspects.

It was estimated that such a huge volume of data were generated due to the following main elements[4-5]:

- Mobility trend: mobile devices, mobile events and sharing, and sensory integration,
- Data access and consumption: Internet, interconnected systems, social networking, convergent interfaces and access models (Internet, search and social networking, and messaging) ,
- Ecosystem capabilities: Major changes in information processing model and the availability of an open framework; the general-purpose computing and unified network integration.

Big Data has characteristics from the aspects of scale, velocity and pattern as the following[1]:

- Means large volume of data in conceptual range as well as in simply stored physical size: Volume of data already exceeded 100 EB at the end of 1990s, reached 1.8 ZB at 2011, and we have already entered in the age of ZB. By 2020, it was forecasted that the volume of data will be 50 times bigger than the one at 2011, which means the main age of ZB.
- Have characteristics that are produced in real-time and are disseminated very rapidly: In 1980s and 1990s, the structured data was a main stream. However, the data became more diverse, complex and socialized in the 2000s and 2010s. In the 2020s and 2030s, the reality and real-timeliness of data will become important characteristics.
- Integrated processing of the existing structured formal data as well as unstructured data uploaded in Internet board, Facebook, SNS etc.. In the 1980s and 1990s, the main stream of data was the structured data such as database, office information, etc. In the 2000s and 2010s, we had entered in the age of unstructured data such as e-mail, multimedia, and SNS. In the 2020s and 2030s, it was

forecasted that we will enter in the age of machine information and cognitive information data such as RFID, sensor and machine-to-machine (M2M).

Three main elements are required for the utilization of Big Data as follows[2]:

- Cloud computing: Because it is difficult to process Big Data with the existing analysis tools, cloud computing technology such as MapReduce, Hadoop, and Hbase is required to analyze and process data.

- Networking environment: In order to implement the analysis result by using the real-time cloud computing technology, the construction of network infrastructure is needed.

- Real-time usability: Regardless where data is generated, it should be possible to use in real-time basis.

As discussed above, due to the exponential growth of digital information volume, Big Data has been an important issue. The government of United States of America had thus established the active utilization strategy of Big Data in March 2012 through the 'Big Data R&D Initiative'[6].

In Korea, communication industry and Big Data have very close relationships. Therefore it is necessary to study domestically on the efficient network infrastructure for Big Data[2]. Accordingly, this paper intends to present the networking considerations and Cisco case study for Big Data.

2. Networking Considerations

2.1 Big Data and Big Traffic

Big Data from multi-site corporate produces big traffic. Big Data applications also induce the massive amounts of traffic with significantly increased real-time and workload-intensive transactions. The movement of large volume of data set over WAN(Wide Area Network) is required to support the Hadoop applications before execution, during execution, and after execution of them. IRG(Internet Research Group) recommends examining on big traffic as earliest as possible when the Hadoop cluster installation is considered and/or planned[7]. The reason is that the scalability and usability of Hadoop cluster may be damaged without understanding the role of WAN in the application of enterprise Hadoop.

The problem of big traffic arises when the processing stages of Big Data and multiple geographically distributed data centers are included. It also happens due to the propagation of data among clusters for the purpose of storage hierarchy management. In these environments, the Hadoop cluster requires a high-speed networking fabric for multi-Giga bits speed. The enterprise networks also should

be optimized to provide a strong infrastructure for the volume, velocity, and accessibility of data, supporting the traditional transaction-oriented RDBMS and various applications such as Big Data[8].

Based on the IDC white paper[8], traffic patterns tends be bursty and variable partly because of the uncertainties of movement of data over the network at any given time. Delays in data transfer were noted to be significant unless the requisite network resources are provided. To achieve appropriate network efficiency, proper line rate performance and rightsizing switch capacity are stated to be necessary.

2.2 Network Characteristics

Typically, one or more of the following phases of MapReduce jobs were found to transfer data over the network, based on [5]:

- 1) Writing data: The initial data is written in HDFS(Hadoop Distributed File System) either by streaming or bulk-delivering. When additional data is transferred over the network, data blocks of the loaded files are replicated.

- 2) Workload execution: The MapReduce algorithm is run in the following four phases:

- Map phase: If the data block is not locally available and has to be requested from another data node (i.e., HDFS locality miss occurs), the network is used at the beginning of the map phase.

- Shuffle phase: In this phase, the intermediate data is transferred between the servers. Data is transferred over the network when the output of the mappers is shuffled to the reducers.

- Reduce phase: In this phase, the data is locally aggregated on the servers. Almost no traffic is sent over the network in this phase because the reducers have all the data they need from the shuffle phase.

- Output replication: MapReduce output is stored as a file in HDFS. The network is used when the blocks of the result file have to be replicated by HDFS for redundancy.

- 3) Reading data: This phase occurs when the final data is read from the HDFS for consumption by the end application, such as the website, indexing, or SQL database.

In addition, it was noted that the network is crucial for the Hadoop control plane: the signaling and operations of HDFS and the MapReduce infrastructure.

Based on the test results, [5] presents the relative importance of parameters to job completion as shown in the following order:

- Availability and resiliency: To provide a network that is available and resilient, the deployed network architecture should provide the required redundancy and that can also

scale as the cluster grows. Switches and routers should also provide availability and resiliency.

- Burst handling and queuing: Because several HDFS operations and phases of MapReduce jobs are bursty, a network is required to handle bursts effectively. Switches and routers with architectures that employ buffer and queuing strategies that can handle bursts effectively should be chosen.

- Oversubscription ratio: Because overprovisioning the network can be costly, it was noted that generally accepted oversubscription ratios are around 4:1 at the server access layer and 2:1 between the access layer and the aggregation layer or core. It was concluded that network architecture that deliver a linear increase in oversubscription with each device failure are better than architectures that degrade dramatically during failures.

- Data node network speed: It was recommended that data nodes should be provisioned with enough bandwidth for efficient job completion, considering the trade-off relationship between price and performance.

- Network latency: It describes that variations in switch and router latency have a minimal impact on cluster performance. A network wide analysis is denoted more important than device level. Moreover, it points out that the latency contribution to the workload is much higher at the application level due to the application logic such as JVM(Java Virtual Machine) software stack, socket-buffer etc than network latency. In any case, it was revealed that slightly more or less network latency will not noticeably affect job completion times.

2.3 Networking Requirements

Due to the change of traffic sources and patterns by the Big Data, network must deal with the phenomenon of traffic shift from server-to-client pattern (which has a traditional enterprise or web server farm characteristics) to server-to-server traffic flow among data center network fabrics. This type of horizontal flow includes links between servers and requires intelligent storage systems to be increased.

Big Data imposes its own computing infrastructure requirements and should incorporate essential functions such as creation, collection, storage and analysis of data. These particular processing requirements are distributed server clusters made up of hundreds or thousands of nodes. In [8], the modular installation of servers at hyperscale has been presented as the preferred method to meet those requirements. The hyperscale server architectures may consist of thousands of nodes which have many processors and disks. Therefore, the networking infrastructure which connects these nodes must be scalable and resilient for the optimal performance, especially when the data are shuffled among them during a certain application phase.

Therefore, it is necessary to have more aggressive and proactive approaches for the planning of Hadoop cluster to support the analysis of Big Data or for the planning of network architecture to support other configurations.

IDC white paper[8] denotes that the network is an essential foundation for transactions between massively parallel servers within Hadoop or other architectures and between the server cluster and existing enterprise storage systems. In [8], the hyperscale network architecture mentioned above is called as "holistic network". The advantages to the holistic network approach were described as the following:

- Ability to minimize duplicative costs whereby one network can support all workloads,
- Multitenancy to consolidate and centralize Big Data projects,
- Ease of network provisioning where sophisticated intelligence is used to manage workloads based on the business priorities,
- Ability to leverage network staffing expertise across the datacenter.

Other factors affecting the design and implementation of networks was noted as governance or regulation requirements[8]. For example, in the application of health care, the separation of data plane may be necessary to meet the privacy requirements of sensitive data in the application.

3. Cisco Case Study

3.1 Unified Network Fabric

The networking considerations for Big Data presented by Cisco are based on the real network traffic patterns of Hadoop framework. By understanding the traffic pattern of an application, the coordination between the application and network design may be possible.

In order to accommodate Big Data, it was proposed in [5,8] that the two main building blocks are being added to the enterprise stack as shown in Fig. 1:

- Hadoop: It is required to provide storage capability through a distributed, shared-nothing file system, and analysis capability through MapReduce.
- NoSQL: It is required to provide the capability to capture, read, and update, in real time, the large influx of unstructured data without schemas (e.g., click streams, social media, log files, event data, mobility trends, and sensor and machine data).

Once the basic enterprise requirements are given, these two Big Data components are integrated with the existing enterprise business model.

Hadoop is necessary to provide the framework to handle massive amounts of data. The purpose is to either transform it to a more usable structure and format or

analyze and extract valuable analytics from it. To efficiently process massive amounts of data, it was noted that it is important to move computing to where the data is, using a distributed file system, rather than a central system, for data. Therefore, Cisco proposes that a single large file is split into blocks, and the blocks are distributed among

the nodes of the Hadoop cluster. In [8], it was noted that this “localized” data/compute model introduces two distinct variables: complex data life-cycle management and matching of nodal capacity in terms of compute and I/O need for variety of workloads.

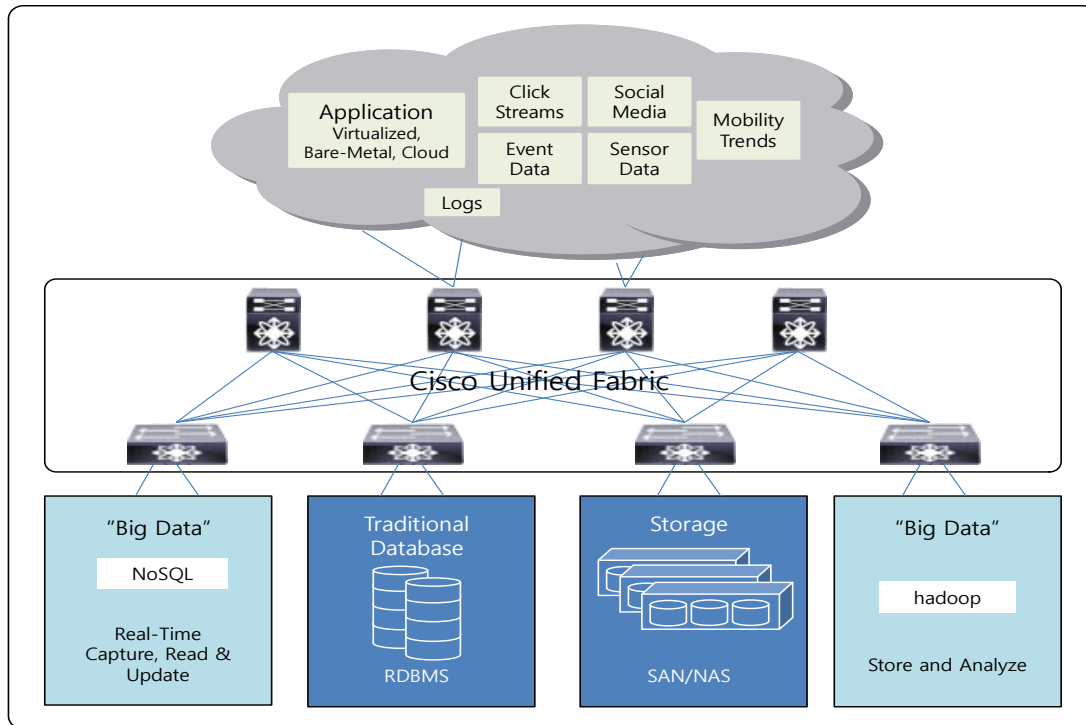


Fig. 1. Big Data Building Blocks and Cisco Unified Fabric[5]

In the IDC white paper[8], this unified Ethernet fabric is denoted as more flatter and converged networks. Through the converged networks, it is stated that the complexity and cost due to the multi-fabric, separate adaptor and cabling may be reduced. It also describes that the flatter-based network architecture maximize the network efficiency, reduces congestion, and may deal with the limitedness of spanning tree by creating the active layer 2 network path for load balancing and redundancy. Compared with the traditional Ethernet fabric, it is mentioned that the unified Ethernet fabric maximizes the performance and availability of applications while the cost and complexity may be reduced.

The unified Ethernet fabric may result in the full link utilization by using the multi-path through the network and by consistently deciding the most efficient path. This architecture also has a superior scalability.

In [8], it was mentioned that the unified fabric brings the following benefits to Big Data:

- Scalability: The fabric can scale incrementally with the growth of big data applications.
- Multitenant architecture: The fabric has ability to provide a multitenant architecture across multiple use cases.
- Machine-to-machine traffic: With resource buffering that is integral to Big Data infrastructure architecture, the fabric was denoted to be designed for machine-to-machine traffic flows

3.2 Test Results

Because many types of workloads can be run to use distributed computing facilities in Hadoop, there are many factors that affect workload completion times. In order to demonstrate the behavior of the workloads in the network, two types of them were used in the test[5].

- Business Intelligence (BI) workload: This workload is a reduced-function workload in which the amount of output data is much smaller than input data. For example, this workload takes 1TB data as an input and outputs 1 MB data.

- Extract, Transform, and Load (ETL) workload: In ETL workloads, a large amount of data needs to be converted to another format suited for various applications. These types of work loads are found most common in enterprises.

When multiple data nodes running mappers finish the map task and reducer pull the data from each node finishing the map task, it is shown that the multiple bursts of receive data exist. It was also found that this traffic is minimal as the data node is performing a compute intensive map task.

From the ETL workload benchmark, the whole event is shown with a data node receiving a large amount of data from all the senders. This is due to the fact that the output of the ETL work load remains the same as input.

In the test result of non-local data impact, it is shown that the initial spike occurs in received traffic before the reducers start. This spike represents data that each map task needs that is not local.

In the test of Hadoop reduce-shuffle phase, it is shown that there is significant amount of traffic because the entire data set needs to be shuffled across the network. The spikes are made up of many short-lived flows from all the nodes in the job and can potentially create temporary burst trigger short-lived buffer and I/O congestion.

In the figure of the aggregate traffic during replication, it also shows the spike caused by multiple nodes sending data at the same time.

4. Conclusions

While the steel and coal, and Internet was the main elements supporting the world economic change during the industry revolution and IT revolution respectively; Big Data is expected to play the main role of economic change during the upcoming mobile smart revolution[3]. In order to use Big Data efficiently, the construction of network infrastructure for the implementation of real-time analysis results by using the Cloud Computing technology is required.

Big Data produces big traffic and thus results in the significant burden to the network infrastructure. Therefore, the enterprise network should be optimized to support a strong foundation in terms of volume, speed, and accessibility of data for both traditional transaction-oriented RDBMS and diverse applications such as Big Data.

In this paper, networking infrastructure considerations to support Big Data are surveyed and Big Data networking architecture is presented through the case study of Cisco.

References

- [1] Korean President's Council on National ICT Strategies, The national development strategy in Big Data age, pp.49-58, August 2012.
- [2] Sung-Choon Lee, "The viewpoint on the Big Data utilization and communication industry", <http://www.ktoa.or.kr>, pp. 6-11, Vol. 60, Spring 2012.
- [3] Sung-Choon Lee, Yang-Soo Lim, Min-Jee Ahn; Big Data: The key to open the future, KT Economics and Management Research Center Report, July 2011.
- [4] Cisco White Paper, Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update 2010-2015, Feb. 2011.
- [5] Cisco White Paper, Big Data in the Enterprise: Network Design Considerations, 2011.
- [6] Eung-Yong Lee, "The USA government Big Data R&D strategy", KISA, Internet and Security Issue, pp.3-26 August, 2012.
- [7] Internet Research Group, Big Data, Big Traffic and the WAN, Jan. 2012.
- [8] Lucinda Borovick and Richard L. Villars, The Critical Role of the Network in Big Data Applications, IDC White Paper, April 2012.



Yong-Hee Jeon received the B.S degree in Electrical Engineering from Korea University in 1978 and the M.S and Ph. D degrees in Computer Engineering from North Carolina State University at Raleigh, NC, USA, in 1989 and 1992, respectively. From 1978 to 1985, he worked at Samsung and KOPEC(Korea Power Engineering Co.). Before joining the faculty at CUD (Catholic University of Daegu) in 1994, he worked at ETRI(Electronics and Telecommunications Research Institute) from 1992 to 1994. Currently, he is a Professor at the School of Information Technology Engineering in the CUD, Gyeongsan, Korea. Since January 2008, he has been a Vice-President of KIISC(Korea Institute of Information Security and Cryptology).