

Chapter 2

Overview of Big–Data–Intensive Storage and Its Technologies

Richard S. Segall

Arkansas State University, USA

Jeffrey S. Cook

Independent Researcher, USA

ABSTRACT

Computing systems are becoming increasingly data-intensive because of the explosion of data and the needs for processing the data, and subsequently storage management is critical to application performance in such data-intensive computing systems. However, if existing resource management frameworks in these systems lack the support for storage management, this would cause unpredictable performance degradations when applications are under input/output (I/O) contention. Storage management of data-intensive systems is a challenge. Big Data plays a most major role in storage systems for data-intensive computing. This chapter deals with these difficulties along with discussion of High Performance Computing (HPC) systems, background for storage systems for data-intensive applications, storage patterns and storage mechanisms for Big Data, and the Top 10 Cloud Storage Systems for data-intensive computing in today's world. Big Data storage server statistics and usage distributions for the Top 500 Supercomputers in the world are also presented graphically and discussed.

INTRODUCTION

Data-intensive computing systems have penetrated every aspect of people's lives. Behind it is the scientific and commercial processing of massive data impacting the decision makings in companies, academics, governments, social cites, and personal lives.

There are two types of data-intensive computing systems that continue to co-exist in the modern computing environment:

DOI: 10.4018/978-1-5225-3142-5.ch002

1. High Performance Computing (HPC) systems, consisting of tightly coupled computer nodes and storage nodes that are used to execute task parallelism for scientific purposes like weather forecasting, physics simulation, and the like (Rouse, 2017b).
2. Message Passing Interface (MPI) is an example of a computing framework on HPC systems. Big Data systems, comprised of more loosely coupled nodes, are used to execute data parallelism for tasks such as sorting, data mining, machine learning, etc. MapReduce is an example of a computing framework on Big Data systems (Barney, 2017).

Both HPC systems and Big Data systems that are deployed for multiple users and applications to share the computing resources so that 1) the resource utilization is high, driving down the usage cost per application/user, and the users get better responsiveness of application execution; 2) the data set is reused without extra overhead to move around performing redundant Input/Outputs (I/O) and users can also save space.

As the computing needs continue to grow in data-intensive computing systems, the shared usage model results in a highly resourceful competing environment. For example, Amazon, Apple and eBay provides HPC and Big Data as cloud services. Hadoop version 2, YARN (Yet Another Resource Negotiator), that is one of the key features in the second-generation Hadoop 2 version of the Apache Software Foundation's open source distributed processing framework. Originally described by Apache as a redesigned resource manager. YARN is now characterized as a large-scale, distributed operating system for Big Data applications which provides a scheduler to incorporate both MapReduce and MPI jobs (Rouse, 2017a).

As the number of concurrent data-intensive applications and the amount of data increase, application I/O's start to saturate the storage and interfere with each other, and storage systems become the bottleneck to application performance. Both HPC and Big Data systems I/O amplification adds to the I/O contention in the storage systems. To counter failures in these distributed systems, HPC systems employ defensive I/O's such as check pointing to restart an application from where it fails, and Big Data systems replicate persistent data by a factor of k , which grows with the scale of the storage system. Both mechanisms aggravate the I/O contention on the storage. The storage systems can be scaled-out, but the compute to storage node ratio is still high, rendering the storage subsystem a highly contended component (Xu, 2016). Therefore, the lack of I/O performance isolation in the data-intensive computing systems causes severe storage interference which compromises the performance target set by other resource managers proposed or implemented in a large body of works. Failure to provide applications with guaranteed performance has consequences. Data-intensive applications must complete in bounded time so as to get meaningful results. For example, weather forecast data is much less useful when the forecasted time has passed. Paid user in a Big Data system also require a predictable runtime even though the job is not time sensitive, and the provider may get penalized in revenues if jobs fail to complete in a timely manner (Xu, 2016).

This chapter addresses the problems stated above for data-intensive computing systems. It provides different approaches for both HPC storage systems and Big Data storage systems because their differences in principles, architecture, and usage pose distinct challenges. Before studying these systems and addressing their respective problems separately, the discussion of the differences between these two types of systems is established here (Xu, 2016).

HPC systems are strongly coupled distributed systems, connected by expensive hardware and network links (e.g. InfiniBand [inf]). The application execution principle focuses on *task parallelism*, and thus both its parallel compute processes and I/O requests are tightly coupled and must be executed *together*.

40 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage:

www.igi-global.com/chapter/overview-of-big-data-intensive-storage-and-its-technologies/198755

Related Content

A Generalized Comparison of Open Source and Commercial Database Management Systems

Theodoros Evdoridis and Theodoros Tzouramanis (2009). *Database Technologies: Concepts, Methodologies, Tools, and Applications* (pp. 13-27).

www.irma-international.org/chapter/generalized-comparison-open-source-commercial/7899

Physical Modeling of Data Warehouses Using UML Component and Deployment Diagrams: Design and Implementation Issues

Sergio Lujan-Mora and Juan Trujillo (2006). *Journal of Database Management* (pp. 12-42).

www.irma-international.org/article/physical-modeling-data-warehouses-using/3351

STEP-NC to Complete Product Development Chain

Xun W. Xu (2006). *Database Modeling for Industrial Data Management: Emerging Technologies and Applications* (pp. 148-184).

www.irma-international.org/chapter/step-complete-product-development-chain/7891

Artificial Intelligence and Machine Learning for Job Automation: A Review and Integration

Gang Peng and Rahul Bhaskar (2023). *Journal of Database Management* (pp. 1-12).

www.irma-international.org/article/artificial-intelligence-and-machine-learning-for-job-automation/318455

A Model for Estimating the Savings from Dimensional vs. Keyword Search

Karen Corral, David Schuff, Robert D. St. Louis and Ozgur Turetken (2009). *Advanced Principles for Improving Database Design, Systems Modeling, and Software Development* (pp. 146-157).

www.irma-international.org/chapter/model-estimating-savings-dimensional-keyword/4297