



Hadoop Business Case: A Cost Effective Queryable Data Archive/Storage Platform

Big Data's open source solution Hadoop is not just a super-sized database. Hadoop's cost per terabyte is much less than high end data warehouse database server like Teradata and Oracle's Exadata. Hadoop's storage costs are also substantially less than many high and mid-level storage area network (SAN) solutions. Since Hadoop is fundamentally a distributed file storage system, it can cost effectively replace other data archiving and/or data storage solutions. Hadoop also has the additional benefit of being able to query its stored data.

Hadoop's basic storage structure is a distributed file system called HDFS (Hadoop Distributed File System). It distributes data very simply but effectively. HDFS automatically makes 3 copies of the entire file across three separate computer nodes of the Hadoop cluster (a node is a commodity Intel server). While duplicating data three times does consume more space than other methodologies, Hadoop's storage cost is so much less than standard systems, it is actually a practical, cost-effective solution. If a node goes offline, HDFS has two other remaining copies. This is somewhat similar to RAID configurations. What is different is that the entire file is duplicated not segments of a hard disk, and the duplications occurs across computer nodes versus multiple hard disks. Please be aware that a Hadoop node can and should be configured for RAID across its internal hard disks. HDFS is not aware of RAID except as a contiguous drive.

A Hadoop cluster is built using racks of commodity Intel servers with each server having internal hard disk storage. One 72 inch floor rack can hold 42 1U servers. While the odds of a server node going offline for any reason is a significantly higher risk than a hard drive failing, HDFS mitigates the risk by making the three copies that provide a high level of overall reliability. Three nodes would need to go offline simultaneously to risk temporary data loss. All three nodes would need to fail catastrophically for permanent data loss. Catastrophic failure is defined as the data on the node's disks being unable for retrieval. Since HDFS duplicates the entire file versus segments of disk, as long as the server node or even just the drive(s) can be brought back online, the file can be retrieved. While a disk failure almost certainly means data loss on that disk, a server can go offline without data loss. Even though Hadoop has a very small risk to data becoming temporary unavailability, it has almost no risk to permanent data loss.

What kind of server is recommended for Hadoop nodes? While you could buy inexpensive desktop class computers to create the Hadoop cluster, this is not recommended. For an enterprise class Hadoop cluster, a mid-range Intel server is recommended. These typically cost \$4,000 to \$6,000 per node with disk capacities between 3TB to 6TB depending desired performance. This means node cost is approximately \$1,000 to \$2,000 per TB.

HDFS has no physical limitations on file sizes. Because Hadoop is also a query engine and HDFS duplicates the entire file, there are some practical limitations for optimal file size. Hadoop does not query extremely small size files well, nor should they be extremely large. By moderating file sizes, Hadoop distributes query loads across nodes and CPUs.

Hadoop has a numerous query tools that expose Hadoop's map-reduce query technology. A detailed discussion of these is beyond the scope of this article. The bottom line is that data stored on Hadoop can be queried by several languages and tools including but is not limited to Java, SQL, R and Ruby-based tools.

Authored by Fred Zimmerman

Fred Zimmerman, a Data Architect for the consulting firm StatSlice Systems (www.statslice.com) is a veteran of data warehouse, business intelligence and database solutions with over 17 years of experience at Fortune 500 companies. Fred has proven experience integrating innovative ideas with industry best practices with the end result being streamlined, scalable and versatile data solutions.

Fred has designed business intelligence solutions for Verizon, Walmart, WellPoint, Coca-Cola Enterprises, Bank One (now Chase), Shell Oil, Microsoft Consulting Services and EDS (now Hewlett Packard).

His numerous specialties include experience with data warehousing, business intelligence and analytics, Master Data Management, all large-scale databases plus Hadoop and Hbase, and most Business Intelligence and ETL platforms (Microsoft, OBIEE, MicroStrategy, and Informatica).

For More Information

For more information about StatSlice Systems products and services, call (214) 206-9290 or email us at info@statslice.com. Please visit us at <http://www.statslice.com>.