Research Report

Veristorm: Hadoop on the Mainframe?

Introduction

Hadoop is an open source database environment that designed to process large volumes of structured and unstructured data (Big Data). It eliminates network and interface bottlenecks and spreads work across large scale computing clusters. It operates on the assumption that hardware fails. Accordingly, when failures occur, the Apache Hadoop framework (common libraries and utilities, a distributed file system, a distributed file system and associated programming model) intervenes to address the failure. From an execution perspective, Hadoop architecture operates using a model that has become known as "single master/multiple worker". And, finally, Hadoop is largely deployed on x86-based server clusters.

> ***Readers should dwell on several of the Hadoop descriptors from the preceding paragraph, particularly: 1) open source database; 2) programming model; 3) addressing failure; 4) master/worker nodes; and especially on: 5) x86-based clusters.***

Now, starting with point #6, and going down this list in descending order, let's reconsider Hadoop architecture:

- *Largely deployed on x86 clusters* – Why? Hadoop is almost always found running on x86 clusters because these servers were the initial development platform for Hadoop; they are low cost; and x86 architecture is heavily used by the open source community. But the big question that needs to be answered is: "Is x86 architecture necessarily the best server architecture for processing large volumes data?" We would argue that in some cases it is not (for instance, when enterprise class security is required).
- *Master/worker nodes* – This mode of workload execution was formerly known as the "master/slave" mode of computing – it had its roots in mainframe computing. Mainframe architecture is still extremely well suited for executing master/slave workloads due to its scalability (mainframe architecture is the most scalable in the industry). So, if the processing task is to manage a substantial number of "worker nodes", the centralized mainframe management design is well suited to this task.
- *Addressing failure* – One of the basic tenets of Hadoop is that systems fail. But, what if systems don't fail? Mainframe reliability is measured in MTBF (meantime-between-failure). Mainframe MTBF is usually measured in decades (20-50 years); x86 MTBF has historically been much lower (see page 8 of this report). Fewer failures lead to more reliable Hadoop processing – and mainframe architecture rarely fails…
- *Programming model* – Hadoop can be programmed using several different programming languages, but JAVA is the most dominant by far. Mainframes actually offer Java program acceleration at the chip instruction set level, leading to superior JAVA performance; and,
- *Open source database* – which means that Hadoop is available to open source communities worldwide provided these organizations live up to pertinent open source license requirements. Hadoop open source code runs very well on the mainframe…

What we hoped readers have taken away from the preceding five bullet points is this: "x86 architecture may not be the best architectural choice when it comes to processing Hadoop workloads." Mainframes offer reliability measured in decades; they also offer JAVA performance acceleration; and thanks to a master/slave architectural design – mainframes are better suited for processing Hadoop workloads than x86-based servers.

In this *Research Report*, Clabby Analytics takes a closer look at running Hadoop on the mainframe. We discuss why the location of Big Data is important in the server decision making process. We examine mainframe analytics processing capabilities. We then discuss mainframe Hadoop Big Data processing environment offered by Veristorm. We conclude that if the Big Data to be analyzed resides with a mainframe domain, it should be processed by the mainframe.

ETL Process and Selecting the Right Server for the Workload

When choosing the best server to execute a big data workload, we suggest that decision makers consider two things: 1) the location of the data; and, 2) the analytics processing characteristics of a given server. When considering the location of the data, we suggest that information technology (IT) buyers become familiar with the consequences of ETLing data (the extract, transform and load process). When considering server processing characteristics, we suggest that IT buyers make their selection based upon how well a given processor/system design can execute a given workload.

The ETL Process

The process of moving data from one platform to another is known as the ETL process. At *Clabby Analytics*, we are also strongly opposed to moving data *unnecessarily* to other platforms – primarily because it calls for additional systems and storage to be purchased – and because data is duplicated many times (causing the potential for data to be mishandled/corrupted).

Consider this typical scenario when data is moved from a mainframe to an x86-based server environment for processing:

1. Mainframe data is extracted, transformed (shipped) and loaded (mounted) on several x86-based servers. This process isn't free – it burns mainframe MIPS, it requires extra software, it requires extra hardware, and it requires a lot of database administrator time and effort to transfer the data, make extra copies, and manage that data;
2. In order to protect this data, database managers usually make one initial copy of this ETL data plus three derivative copies. Think of the opportunity for human error or machine error in this scenario if the data is mismanaged or corrupted;
3. The ETL process requires an enterprise to invest in even more storage in order to handle the copy and derivatives transferred;
4. There are also additional costs for networking/communications equipment;
5. Additional systems are required to process the ETLed data; and,
6. Expensive human resources (database managers) are required to manage the ETL process and related equipment.

When selecting the right server to process a given workload, we always first ask: "where is the data located?" If the server that owns the data can execute the given workload efficiently, we almost always recommend that this server be selected to process analytics workloads in order to avoid additional ETL costs. For more on this discussion, see our report entitled "The ETL Problem".

Server Selection

At *Clabby Analytics*, we are strong believers in the concept of matching workloads to the servers best suited to execute those workloads in order to achieve optimal performance. Enterprises that execute their workloads on servers that are optimized to process those workloads consistently see faster results and higher return on investment (ROI) than those that do not.

> ***When it comes to selecting the right server environment, the primary goal for any information system selected should be to achieve balanced performance between processors, memory, and I/O (input/output). Accordingly, we recommend that IT executives try to avoid scenarios where a workload overpowers the processor, memory or the I/O subsystem.***
>
> ***To avoid scenarios like this, we recommend that IT executives look at CPU performance characteristics, off-load characteristics (such as the availability of accelerator systems), memory utilization characteristics, processor core efficiency, execution styles, instruction set characteristics, and communications and network facilities.***
>
> ***We also recommend that the system design be scrutinized. We look for Quality-of-Service extensions that have been designed into the system (such as redundant components for availability, or autonomic management environments that can predict failures before they occur)? What are the power consumption characteristics? Is the system design purpose-optimized?***
>
> ***The following section takes a closer look at IBM's System z mainframe from both a processor and system design perspective.***

Is the Mainframe Well-Suited to Handle Analytics Workloads?

There are IT executives in this industry who will argue that the mainframe is not well suited for processing analytics workloads – instead, they argue, the mainframe is only good at processing large volumes of transactions. In a way these executives are correct; in a way they are not.

Where They Are Right

One of the most important aspects to understand about the mainframe processor is that it is a high volume, single-thread processor that accepts stacked work. What this means is that large volumes of work can be stacked-up and processed by the fastest processor in the industry – the z processor. This is why the mainframe excels at transaction and batch processing – work gets stacked and executed rapidly using the z processor. Both x86 and POWER8 processors are multicore/multi-threaded processors that lend themselves nicely to processing parallel workloads (many analytics workloads are parallel-oriented). So these processors can execute certain analytics workloads that exploit parallelism better than a z processor. (Note: Hadoop can process batch workloads – so batch-oriented Hadoop workloads on a System z make perfect sense).

Where They Are Wrong

The System z processor can now rival the performance of distributed systems processors when it comes to mathematics-intensive processing. This has not always been the case, which is why ETL processes were originally put in place. We took a close look at the z processor in this microprocessor report and found that: "IBM introduced IEEE binary floating point facilities [better mathematical processing facilities] at the end of the 1990s. The early 2000s brought 64-bit computing and superscalar parallelism to the mainframe (superscalar architecture implements a form of parallelism called "instruction level parallelism" ─ allowing a single processor to process

work at a rate faster than its clock rate). Also, clock speeds have continually increased. Further, IBM recently added "out-of-order execution", and has substantially improved floating point performance (mainframe floating point now rivals reduced instruction set processors such as POWER, SPARC, and Itanium)". It also rivals x86 mathematical processing capabilities.

Further, System z System design has improved, making it more suitable for analytics processing. Analytics activities tend to require a lot of memory and cache, as well as a speedy input/output (I/O) subsystem to expeditiously feed data to the processors. IBM's mainframe offers more three terabytes of main memory; significantly more cache than its predecesor; and improved I/O speed.

Additionally, more analytics software is available on System z than ever before. About five years ago IBM executive management mandated that the company's software products are to be developed to work across all IBM servers (this includes System x, Power Systems, and x86 servers). In days gone by IBM's Cognos, SPSS and other analytics products may have favored distributed systems – but this is no longer the case. So the same high quality reporting tools, utilities and applications that were perfected and honed to perform analytics on distributed systems are now available and have been optimized for use on the mainframe. And add to this picture that IBM has spent $24 billion on analytics software acquisition and development – far, far more than other systems competitors – and this level of investment shows up in IBM broad analytics portfolio.

Finally, IBM offers an accelerator to accelerate the processing of complex analytics workloads. Two years ago *Clabby Analytics* started reporting on the evolution of a new approach to performance acceleration: the use of server accelerators. The first accelerator that we described was a tightly attached mainframe appliance known as the IBM DB2 Analytics Accelerator (see this report). What this accelerator does is it receives a snapshot of an IBM DB2 database (which is regularly updated) – and it performs complex analytics on that database at speeds exponentially faster than traditional general purpose server environments. IBM's DB2 Analytics Accelerator has been very well received by the mainframe community.

> ***This accelerator is extremely important to the mainframe analytics story. It acts as a tightly coupled side-car that can be used to offload analytics processing that runs best on parallel systems architecture (this architecture is a fast, field programmable gate array/Intel processing system). It enables mainframe data to stay where it belongs – in the mainframe environment – while off-loading complex query processing to a tightly-coupled mainframe-attached server environment.***

How Does Veristorm Fit into This Picture

After explaining the ETL problem – and after explaining mainframe analytics processing characteristics, some IT executives are willing to concede that keeping data under the control of the mainframe (if the data already resides there) is a good business decision because it eliminates the need to ETL data. Some also accept the idea that the mainframe can serve as highly-scalable, very secure analytics processing environment. If you are one of those executives, read on…

Veristorm – The Maker of Hadoop Big Data Analytical Solutions for the Mainframe

Veristorm, Inc. is a maker of analytics and performance management products for mainframe environments. Its products include zDoop, vStorm Performance Manager and vStorm Enterprise. Veristorm's zDoop offering is the first commercial implementation of Hadoop on the mainframe. It runs on the Linux operating environment as part of Veristorm's vStorm Enterprise environment (see Figure 1).

Figure 1 – Veristorm's zDoop and vStorm Environments

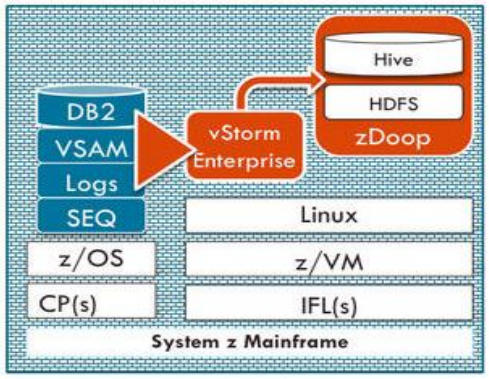

Source: Veristorm – October, 2014

A closer look at Figure 1 shows DB2, VSAM, log and/or SEQ flat files moving from storage under the z/OS operating environment through vStorm Enterprise to zDoop – where those files are stored as Hive or HDFS (Hadoop distributed file system) files running in mainframe Linux partitions. (More precisely, these files are running on the System z "Integrated Facility for Linux" [IFLs] under the control of IBM's z/VM [virtual machine] operating environment upon which Linux runs as a guest operating environment). IFLs are essentially dedicated Linux processor environments that cost less to operate than standard z processors running the z/OS environment. Unlike the traditional ETL process, mainframe z/OS MIPS (millions of instructions per second – the way that mainframe usage is charged) are not being used as data is moved to Linux partitions – saving enterprises processing costs that are associated with external ETL data transfers.

> ***Veristorm has a name for its internal Hadoop transfer from z/OS to Linux control. They call this process the "extract-Hadoop-transform" (EHT) function. This process costs far less than traditional ETL data transfers to external systems – and it also enables data managers to obtain results in near real time.***

According to Veristorm, the primary benefit of using its EHT process is that it simplifies data management. As Veristorm puts it: "It alleviates the need for SQL extracts, ETL consulting engagements with simple point and click data movement from z/OS to HDFS. EBCDIC code page and BCD issues are handled in-flight without MIPS costs, and users avoid the complexity of COBOL copybooks and the security risks, compliance challenges and delays of offloading data".

Also important to note, Veristorm's EHT process is accomplished using a graphical user interface that supports Hadoop, Java, Linux and drag-and-drop management for z/OS data access. In other words, its straightforward to use.

Finally, it should be noted that vStorm Enterprise can send data to external sources for processing. Using vStorm Enterprise, distributed Hadoop configurations can work with Cloudera, Hortonworks, and BigInsights platforms.

> ***Veristorm should be used when data security is of paramount importance (because Veristorm runs on the mainframe and is thus protected by the most powerful security system in the market); and Veristorm should be used to process Hadoop workloads that can run well on a single-threaded stacked processor (the z processor). So batch Hadoop applications clearly belong on the mainframe. Other more complex analytics workloads should be run on IBM's DB2 Analytics Accelerator.***

Veristorm Discussed at the IBM SHARE Conference

Bi-yearly, an IBM system user organization known as SHARE meets to learn more about IBM products and strategies (SHARE is attended by a thousand or so IT professionals – many of them are practitioners looking to learn more about the products that they use and to exchange information with fellow users). One session at this year's SHARE conference featured Dr. Cameron Seay, Ph.D, from North Carolina Agricultural and Technical State University – and he spoke about his experience with Hadoop on the mainframe using Veristorm's zDoop and vStorm Enterprise product offerings. A description of Dr. Seay's presentation can be found here.

In short, Dr. Seay believes that Hadoop is a "game changer" because it eliminates interface and network traffic jam and represents a new way to move data such that it can be processed quickly (it dives a job across many computers – making them more productive). By using Hadoop, Seay believes that Big Data resident in transactions, logs and found in machine data can be processed more quickly – including structured transactions as well as several different types of unstructured data such as data found in system logs, machine or sensor data, emails, images, video, audio and more.

If this data resides in mainframe databases, Seay observed, it must be translated from a proprietary mainframe database to a format more familiar to Big Data analytics. And this is where products such as Veristorm's zDoop and vStorm come in. With these products, Seay pointed out, traditional mainframe operation applications (such as transaction processing) can transfer data to analytical applications. And, as we pointed out earlier, this transfer of data within the mainframe avoids the ETL process and enables analytical insights based on operational data to be found in near real time.

As a user of Veristorm products, Seay described how vStorm Enterprise helps improve mainframe security because "data never leaves the box" – and because the data can take advantage of hardware encrypted streaming and IBM's RACF security environment. And Seay also emphasized that vStorm Enterprise and zDoop are easy to use given the products' graphical user interface, the fact that no programming is required, conversions are automatic, and there is no disk staging required.

As for performance, Dr. Seay shared the following slide (Figure 2) with the SHARE audience. In his tests, he and his students were able to transfer 2 billion records to a Hadoop database and analyze that database in two hours. The data was processed using only 2 Integrated Facility for Linux partitions. This provides a strong argument and proof point that System z is a very powerful analytics processing environment.

Figure 2 – Veristorm zDoop Performance Characteristics

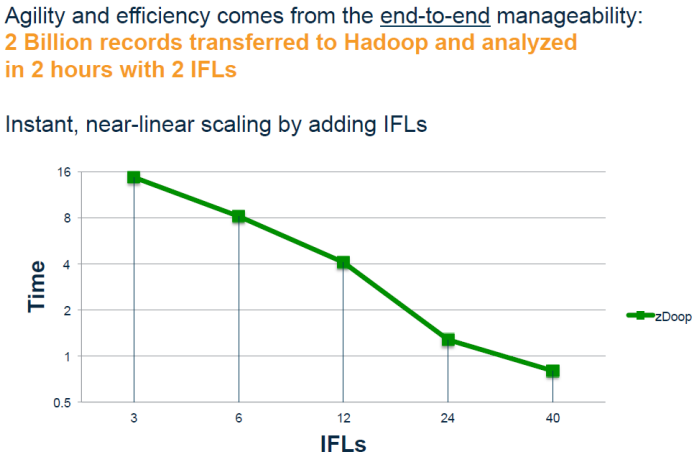

Source: Veristorm – August, 2014

Veristorm: Hadoop on the Mainframe?

Summary Observations

Perhaps the biggest challenges for Veristorm are to convince IT organizations that, for certain workloads (such as batch-oriented Hadoop workloads) they no longer need to ETL data to distributed data marts and data warehouse because the mainframe excels at processing these types of workloads. With Veristorm products, mainframe data can now be shipped to a different region on the mainframe (the Integrated Facility for Linux) region where it can be analyzed.

Last year *Clabby Analytics* wrote this microprocessor report in which we pointed out that over the past several years IBM has been working diligently to improve the compute facilities in mainframe hardware. In this report we stated that: "IBM introduced IEEE binary floating point facilities [better mathematical processing facilities] at the end of the 1990s. The early 2000s brought 64-bit computing and superscalar parallelism to the mainframe (superscalar architecture implements a form of parallelism called "instruction level parallelism" ─ allowing a single processor to process work at a rate faster than its clock rate). Also, clock speeds have continually increased. Further,IBM recently added "out-of-order execution", and has substantially improved floating point performance (mainframe floating point now rivals reduced instruction set processors such as POWER, SPARC, and Itanium)". The mainframe processor is now more capable than ever before at processing analytics workloads. But note, highly parallel-oriented workloads should be handed off to the tightly coupled DB2 Analytics Accelerator environment.

> ***Note that some analytics workloads are "data-intensive" and are better served by a multi-core, multi-threaded massively parallel architecture (the mainframe is a single threaded stack processing environment). To make the mainframe a good data-intensive processing environment, three years ago IBM introduced its IBM DB2 Analytics Accelerator – a massively parallel architecture that can process data-intensive workloads up to 100X faster than traditional general purpose server designs. This sever is tightly coupled with the mainframe; data is snapshotted to it and refreshed constantly – and its performance rivals that of competing servers such as IBM Power Systems, Oracle SPARC and Hewlett-Packard Itanium-based servers.***

In short, the mainframe combined with its accelerator is an excellent analytics processing environment. By using Veristorm's zStorm Enterprise and zDoop offerings, the mainframe can now be used as a powerful Hadoop processing engine for workloads that can be stacked and executed on the System z single-thread process. Other more complex workloads should be handed over to IBM's DB2 Analytics Accelerator.

In the end, it all comes down to this – if the data resides in mainframe databases, that data should be processed by the mainframe. Veristorm's extract-Hadoop-transform approach enables certain stack-oriented Hadoop workloads to be processed on the mainframe – offering enterprises a way to efficiently process structured and unstructured data on the most secure system in the industry.

Clabby Analytics
http://www.clabbyanalytics.com
Telephone: 001 (207) 846-6662