THE CLIPPER GROUP
Navigator™
Navigating Information Technology Horizons℠
Published Since 1993 Report #TCG2016009 October 21, 2016

IBM z Systems Get a Big Charge Out of Spark Analytics

Analyst: Stephen D. Bartlett

Management Summary

About 1 billion years ago, give or take a couple of 100 million, the first recognizable plant life began its emergence on planet Earth. Today, even cursory observation reveals an amazing and beautiful array of flora, and all that grows from the lowly lichen to the mighty cypress. That's the miracle of evolution. The first animals emerged from the sea about 400 million years ago; mankind began its emergence out of the African continent about 50 thousand years ago. Soon Earth's population will reach 7.5 billion, organized in a myriad of organized communities, speaking and writing in thousands of languages and dialects. That's another miracle of evolution.

Fifty-two years ago, *IBM* announced the *System/360*, the first family of unitary architecture computers that changed the nature of information technology profoundly[1]. More recently, despite the many predictions of the IBM mainframe's extinction and irrelevance, that family continues to prevail and evolve; still the mainstay of thousands of enterprises throughout the world. Its most current manifestations are the *z Systems z13*[2] and *z13s*[3], and the *LinuxONE Emperor* and *Rockhopper*[4]. *What's the secret of its longevity?* Evolution, of course. Underneath the covers of today's systems, the architecture of the original System/360 can still be found, but now evolved to meet the challenges of the modern world. *What does that world look like?*

Without dilution of the highest standards for the qualities of service exemplified in IBM's mainframes very early in its life – availability, recoverability, scalability, and security – these world-leading enterprises have established the mainframe as their *systems of record* and the mainframe now is taking leadership positions as the co-located *systems of engagement* and *systems of insight* in a market that eschews proprietary technology and embraces open systems models. Today, the required scope is global, mobile, always on, with demand for instantaneous response in an environment that more closely resembles chaos rather than ordered discipline. Data is streaming in at volumes that are growing exponentially; the challenge is to not only store and

IN THIS ISSUE

[1] See **The Clipper Group Navigator** dated April 8, 2014, entitled *IBM zEnterprise is Enterprise Cloud Infrastructure*, which is available at http://www.clipper.com/research/TCG2014008.pdf.

[2] See **The Clipper Group Navigator** dated January 14, 2015, entitled *The IBM z Systems and the New IBM z13 – Ready to Transform Your Enterprise*, which is available at http://www.clipper.com/research/TCG2015001.pdf.

[3] See **The Clipper Group Navigator** dated February 16, 2016, entitled *IBM z Systems Opens Up Secure Clouds and Introduces the z13s*, which is available at http://www.clipper.com/research/TCG2016003.pdf.

[4] See **The Clipper Group Navigator** dated February 16, 2016, entitled *Open for the Digital Business – IBM Announces LinuxONE Systems*, which is available at http://www.clipper.com/research/TCG2016002.pdf.

manage them, but also to unlock the value contained therein. The Clipper Group has taken the position that there are numerous advantages to co-locating your *systems of record* with your *systems of insight*[5] and that is the first premise guiding this paper.

Additionally, IBM has recognized that participation and leadership in the existing market demands it must be active in an open systems approaches for its solutions that aren't a part of the unique technology of IBM's mainframes. Recently, *Apache Spark*, an open systems analytics environment, has been enabled for native execution on z Systems' *z/OS* operating system. The focus of this paper is to describe this implementation and the many advantages, including economies that it delivers to many enterprises. Please read on if want to know how it may benefit you and your enterprise.

The z Systems Transformation Continues

Mainframe watchers or not, the themes that are dominating the industry today are familiar to all IT professionals. Most prominent are:

- **The development cloud architectures**, particularly hybrid models;
- **Analytic computing**, with focus on the need for real-time response;
- **Support of mobile end-points**, with its diversity of type and its users;
- **Social media data flow**, its volume, variety, and velocity; and
- **Security** in a world inhabited by professional data thieves and malicious players.

These are the primary initiatives driving the current transformation of z Systems development and evolution that are demonstrating the mainframe's relevance in the current environment. These are not in exclusion, of course, of the dominant role recognized as the enterprise system of record that largely is centered on secure transaction processing and recordkeeping.

z Systems as the System of Insight

In order to position the mainframe's role as an enterprise *system of insight*, it is useful to review briefly the foundation that has already been laid in this arena. The central argument for a mainframe-based system of insight is the many benefits of co-location with the *systems of record*, i.e., the benefits of running on the same server and sharing access to the same real-time data. It takes a lot of processing time, memory, and network bandwidth (not to mention absolute time) to extract the data from the system of record and push it over the network to other (outboard) computer systems. So even before the exported data is made available for analytical insights, it has cost a lot of money and time, during which the data often becomes increasingly stale. The other half of the story is all of that data that has been exported has to be stored, transformed into something usable and then stored again, possibly many times, including but not limited to backups. This is a never-ending cycle of costs and time – all just to end up with data of varying (and potentially unknown) degrees of staleness.

The obvious question is *why not operate on the enterprise data at the point where it is permanently recorded*, i.e., run the systems insight processes and applications on the always-up-to-date data sitting in the system of record? Several reasons have been proffered not to consider doing this: fear of somehow affecting or contaminating the processing of mission-critical business, expending the seemingly most expensive computing resources to do what seems to be straightforward work that could be done on less-capable but presumably less expensive servers, and the lack of robust middleware and tools to accomplish this. It easily can be demonstrated that these issues have been addressed through the robust z Systems virtualization capabilities and the advanced monitoring and management facilities of the mainframe platform.

There are several examples of the power of z Systems as the primary platform for analytics.

- ***IBM DB2 Analytics Accelerator (IDAA V5.1)*** accelerates in-database transformations and analytics. The *Accelerator*, as it is called, is a logical extension of DB2 that does not require any changes to existing applications programs. The Accelerator complements the industry-leading transactional processing capabilities of *DB2 for z/OS*, providing a specialized access path for data intensive queries. This enables real and near real-time analytics processing,

[5] See **The Clipper Group Navigator** dated December 21, 2012, entitled *Addressing New Business Analytics Challenges – When the IBM zEnterprise Really Makes Sense*, which is available at http://www.clipper.com/research/TCG2012030.pdf and a companion paper by my colleague Mike Kahn, **The Clipper Group Navigator** dated February 17, 2015, entitled *Why the IBM Mainframe is the Right Place for Enterprise Systems of Engagement and Insight*, which is available at http://www.clipper.com/research/TCG2015002.pdf

which executes transparently to applications and users.

- ***IBM DB2 10.5 with BLU Acceleration*** is available on *Linux on z Systems*. It enables faster queries on Linux data marts without time-consuming custom optimization by compressing row-oriented tables and then operating on the in-memory compressed data.
- ***IBM InfoSphere z Systems Connector for Hadoop*** enables efficient sharing of mainframe data with *IBM InfoSphere BigInsights*, running on Linux for z Systems partitions[6]. When both IBM InfoSphere z Systems Connector for Hadoop and IBM InfoSphere BigInsights are installed on the mainframe, the z Systems platform behaves, in effect, as a large, secure private cloud, greatly extending the range of data that is accessible to the mainframe.
- ***IBM InfoSphere Information Server*** – z Systems' business analytics solution structure also is supported by the IBM InfoSphere Information Server. This integrated software suite provides the essential functions of information integration and governance including ensuring data quality, providing a single view of Master Data, and performing lifecycle management, privacy and security functions that are essential to a complete business analytics solution. The suite is supported on Linux for z Systems.[7]
- ***IBM SPSS*** is another essential analytical tool that is available on the mainframe. Its product portfolio includes world-class facilities for statistics, modeling, data collection, and deployment. The statistics family of tools is the most widely used suite of statistical software in the world and includes facilities for linear and non-linear models, simulations, customizing tables, data preparation, and validity checking. The modeling family of tools includes accessing operational data from a variety of sources, such as Cognos BI, DB2, IBM *Netezza (IDAA)*, Microsoft *SQL Server*, and Oracle *MySQL*. Scoring functions have been ported from *SPSS Modeler 15 to DB2 for z/OS*, thereby leveraging the z Systems' very high speed and reduced latency by eliminating ETL processes and the infrastructure that supports it.

[6] Also available on IBM Power Systems clusters or Intel x86 architecture systems.

[7] See footnote 6 above.

Opening the Door to Open Source Analytics

IBM, despite its erroneous reputation for proprietary (and therefore exclusionary) mainframe systems, has been an active participant in the open systems movement going back to 1988 and the founding of the *Open Software Foundation*, along with several other industry leaders.[8] In the late 1990's this group was greatly expanded to form *The Open Group*, the certifying body for the UNIX trademark.[9] The *UNIX System Services* component of z/OS is tightly integrated into the operating system and remains a key element of IBM's open and distributed computing strategy. *WebSphere Application Server, CICS, IMS, Java Runtime, Tuxedo*[10], *DB2, WebSphere MQ, SAP ERP*[11], *Lotus Domino*, and *Oracle Web Server* all use z/OS UNIX System Services.

Then, in 1999, IBM made the revolutionary introduction of *Linux on System z*[12]. Though it seemed to be a modest nod toward open systems technologies, it proved to be among the most profound innovations in mainframe technology since the introduction of the unified *System/360* architecture. The following year was marked by the establishment of *The Linux Foundation*, a consortium dedicated to fostering the growth of Linux and, as a consequence, Linux on z Systems. Within the foundation's purview, IBM and a number of academic, government, and corporate partners, instituted the *Open Mainframe Project.* The project is a direct response to the critical mass of users and vendors that needed to work together to advance Linux tools and technologies, the goal of which is to foster the increase of enterprise-scale innovation. A key part of IBM's code contributions embodied a range of predictive analytics solutions that constantly monitor for unusual system behavior, constantly identifying and isolating issues that have the potential of turning into failures. This code is in the form of the *zAware*[13] capabilities that were announced for

[8] Several of whom only now exist in the bones of their successors, including DEC, Sun, and Apollo.

[9] IBM's variant is known as AIX. It appeared on the mainframe in 1988 as AIX/370, later AIX/ESA.

[10] *Tuxedo*, Oracle's transaction processing system for UNIX.

[11] *SAP ERP*, formerly *SAP R/3*.

[12] Now called Linux on z Systems.

[13] For more detail on zAware, see **The Clipper Group Navigators** entitled *The IBM zEnterprise EC12 - Bigger, Better, Faster*, dated August 28, 2012, and available at http://www.clipper.com/research/TCG2012019.pdf and *Addressing New Business Analytics Challenges - When the IBM zEnterprise Really Makes Sense*, dated

z/OS in August 2012 and was, at that time, the largest single contribution from IBM to the open source community. This important tool has been ported by IBM to Linux on z Systems and substantially furthers the goal of providing better resiliency and manageability to mainframe-based open system environments.

Data Drives an Important Open Source Project – Apache Spark

IBM has been a member of the *Apache Software Foundation* since its founding in 1999 with the mission to provide organizational, legal, and financial support for the *Apache HTTP Server Project*, the most dominant open source Web server world-wide and, later, other Apache open-source software projects, now numbered in the hundreds. The initial release on the mainframe of *Apache Spark* in May 2014 was in response to the complex analytic needs for "big data" and the limitations of existing paradigms. In June 2015, IBM announced a major commitment to advance Apache Spark with plans to embed Spark into its industry-leading analytics and commerce platforms and to offer Spark as a cloud service on *IBM Bluemix*. This commitment also includes open sourcing its breakthrough *IBM SystemML* machine learning technology that originated in IBM Research laboratories.

In addition, IBM is committing more than 3500 researchers and developers to work on Spark-related projects at more than a dozen labs worldwide and has opened a Spark Technology Center in San Francisco for the data science and developer communities to foster design-led innovation in intelligent applications. At the time, Apache Spark was supported for Linux on z Systems.[14] IBM doubled down in March 2016 by announcing the support of Apache Spark *natively* on its mainframes (i.e., on z/OS) via the *IBM z/OS Platform for Apache Spark.* This implementation enables data scientists to analyze data in place on the system origin, the *systems of record*, without the need to extract, transform, and load (ETL), by breaking the tie between the analytic library and the underlying file system(s). The new platform helps enterprises derive insights more efficiently and securely in many ways.

- **Streamlined development** – Developers and data scientists can use their existing expertise with programming languages such as *Scala, Python, R, SQL, and Java*, to reduce time-to-value for actionable insights.
- **Simplified data access –** Optimized data abstraction services remove complexity, providing seamless access to enterprise data in traditional formats such as *IMS*, *VSAM*, DB2 for z/OS, partitioned data sets (PDSE) and *System Management Facility* (SMF) records via familiar tools and Apache Spark APIs.
- **In-place data analytics** – Apache Spark uses an in-memory approach for processing data to deliver results quickly. The platform includes data abstraction and integration services that enable z/OS analytics applications to leverage standard Spark APIs (application program interface). This allows the organization to analyze data in-place, avoiding the costly processing, complexity, and security considerations associated with of extraction, transformation, and load (ETL) workloads.
- **Open source capabilities –** The platform offers an Apache Spark distribution of the open source, in-memory processing engine that is designed for big data.

An "Operating System" for Analytics

Apache Spark is much more than an analytic engine; it is better perceived as an operational framework that serves multiple constituencies that enables highly iterative analysis on large volumes of data.

- ***The data scientist*** seeks to discover new, actionable insights by identifying patterns, trends, and opportunities. Spark supports the entire data science workflow, from data access and integration, to machine learning and visualization. The data scientists are the analytical gurus that have little concern for the specifics of the operational platform.
- ***The information management specialist/ application developer*** builds the application that use advanced analytics in partnership with the other constituencies. Their tools are the leading analytic programming languages, e.g., SQL, Java, R, Python, and Scala[15]. They build libraries that simplify development operations (DevOps) and minimize complexity.
- ***The data engineer*** is the bridge between the applications developers and the data scientists. They are the closest to the platform and are responsible for putting the right data system

December 21, 2012, and available at http://www.clipper.com/research/TCG2012030.pdf.

[14] Also available on AIX and Linux on IBM Power Systems.

[15] Scala is an object-oriented, functional programming language similar in syntax to the C language and is interoperable to Java.

Exhibit 1 — Where Apache Spark Fits within the Mainframe Architecture

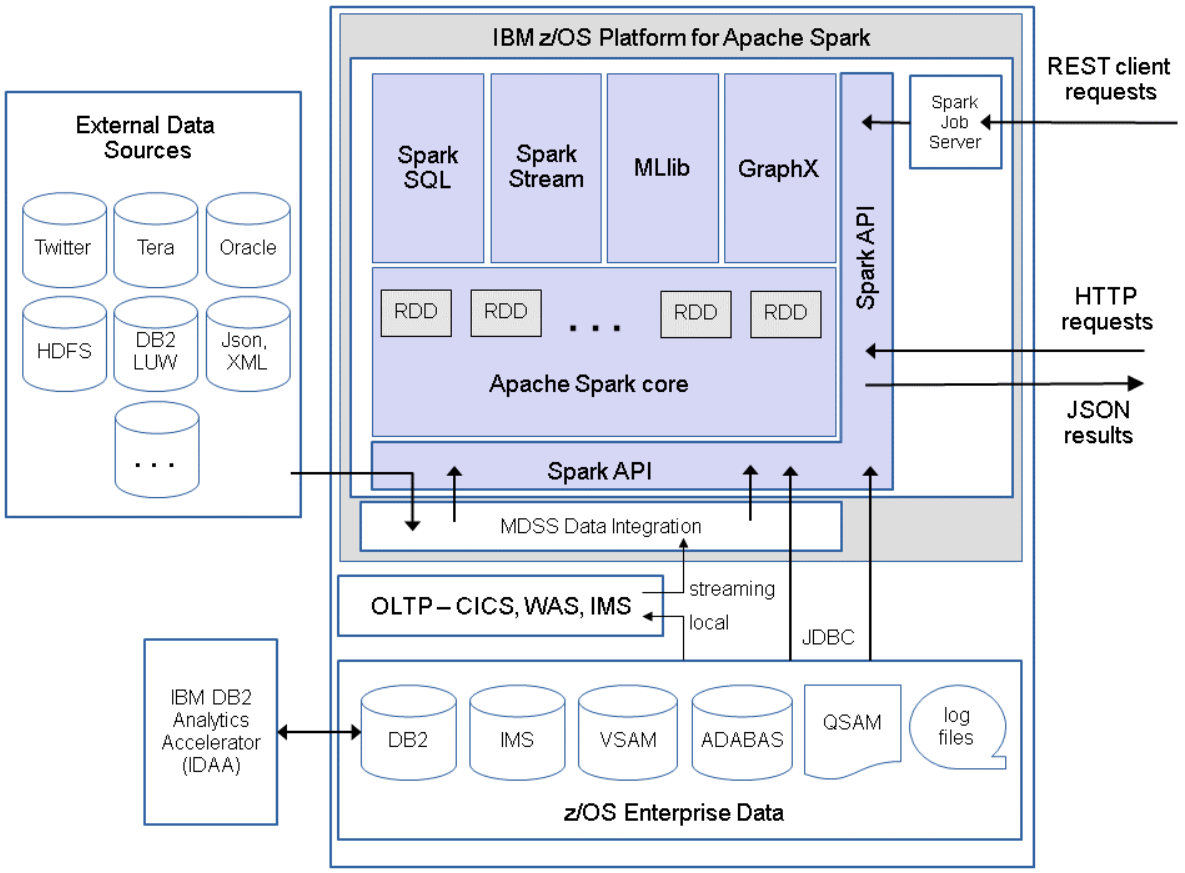

Source: IBM

framework together for the current tasks. Because Spark is data-store agnostic, the data engineers' role is to abstract the complexity of data access from multiple sources. These roles are not unique and frequently overlap but the architecture of Apache Spark provides the common framework for all.

Apache Spark on the z/OS Platform

The architecture of the Spark environment on a z/OS platform is highly flexible and can accommodate the needs of many different enterprise analytic configurations. On the z System-z/OS server side, Spark is deployed with its associated subsystems and components. The off-platform environment provides the user interfaces that are primarily used by the data scientists, information managers, and application developers. (See Exhibit 1 above.)

The Apache Spark server consists of several key components: a core set of services, a set of libraries the enable key functionality, and an API. (Optionally, there is a *Spark Job Server* that enables access to Spark services through a RESTful API to enable access through a web browser.)

In addition to the Spark Core that provides distributed task dispatching, i.e., scheduling and basic I/O functions that are accessed through the Spark API, there are several key components that are essential to the Spark environment.

- ***Spark SQL*** allows developers to intermix SQL with Spark's programming language, supported by Java, Python, and Scala.
- ***Spark Stream*** enables processing of live steams of data, such as log files and queues of messages from sources such as *Twitter* and TCP/IP sockets.
- ***GraphX*** is a graph-processing library with APIs to manipulate graphs and perform graph-parallel computations.
- ***MLib*** is the machine learning library that provides a number of ML algorithm types. IBM donated *SystemML* from its Research Division laboratories as open source to provide additional and more robust machine learning capabilities.

Resilient Distributed Datasets (RDD)

An RDD is the fundamental data structure of Spark. It is an immutable, fault-tolerant collection of elements that can be operated on in parallel across a cluster; it uses caching, persistence (memory, spilling, and disk), and check-pointing. Many database or file types are supported. An RDD is physically distributed across the cluster, but is manipulated as one logical entity. Spark distributes any required processing to all partitions where the RDD exists and performs necessary redistributions and aggregations.

Mainframe Data Services

Unique to this implementation of Apache Spark is an optimized data integration layer called *Mainframe Data Services for Apache Spark for z/OS (MDSS)*. It enables data from multiple z/OS sources, such as DB2 for z/OS[16], *IMS/DB, Adabas, VSAM, QSAM,* physical sequential data sets, PDSEs, SMF, and *SYSLOG*.[17] Online transaction processing (OLTP) can be performed within this environment using *IBM Customer Information Control System (CICS), IBM WebSphere Application Server (WAS)* or *IBM Information Management System (IMS)*, although MDSS is not intended to support low-latency, real-time transactions. It also provides the facilities to integrate off-platform external data sources, such as *Distributed Resource Data Access (DRDA)* data sources within *DB2 for Linux, UNIX* and *Windows* (DB2 LUW), Oracle, Json XML, and the *Hadoop Data File System* (HDFS). With this extensive array of supported data sources and analytical capabilities, Apache Spark for z/OS demonstrably is an ideal platform for unified data integration in support of z Systems as a *system of insights*.

Apache Spark – MapReduce Comparison

Much has been written and claimed for *Hadoop MapReduce* as the primary means for driving "Big Data" analytics. This architecture falls far short when compared to Apache Spark. Hadoop's MapReduce is disk-dependent, with much redundancy; Spark relies on in-memory capabilities which are particularly well-supported on z Systems, resulting in near real-time capabilities. Hadoop MapReduce is a batch-only system whereas Spark supports batch, streaming and interactive modes. Hadoop MapReduce language support is only Java; Spark supports Java, Python, Scala, and R. Spark includes integrated modules for streaming, SQL, graph, and machine learning; Hadoop MapReduce as a loose ecosystem of various tools that must be integrated separately. Apache Spark is superior in its use of resources, execution speed, and its efficiency, and if your enterprise has already implemented an HDFS data store, Spark easily can accommodate that data source as well.

What z Systems and z/OS Brings to the Game

The advantages of co-location of data and the data analysis platform cannot be overstated. Aligning z Systems and z/OS, the enterprise platform of the *system of record*, with the *system of insight*, embodied by Apache Spark for z/OS, leverages the unique capabilities of IBM's mainframe and architecture. Doing this has the potential to drive new capabilities and opportunities for enterprises to expand their application portfolios deeper and wider in markets that previously were not technically or economically feasible.

Complex Analytics on a Single Unified Platform

Apache Spark provides real-time analytics of mission-critical operational data and machine learning that can actively enhance its statistical algorithms. It supports SQL standards (*SQL 92* and *SQL 99*) that are much richer than standard Spark SQL. Access to many different data sources, particularly those native to mainframe systems, is supported as well as to IBM's primary data management systems, DB2 and IMS. In addition, many non-mainframe data sources, such as NoSQL, XML, unstructured data streams, and HDFS, are supported and integrated equally without the need for complex and expensive ETL processes.

Data Security Managed by the Most Secure Platform

Moving your enterprise's proprietary and most sensitive information to a separate analytical platform engenders risks that are difficult and expensive to manage and a monstrous problem to fix, if ever broken. Other issues such as multiple unsynchronized copies, latency, and ownership must also be addressed as well.) The IBM z/OS platform for Apache Spark provides advanced security for SQL, NoSQL, events, and services solutions typically configured by system

[16] And by a transparent extension, the IBM DB2 Analytics Accelerator (IDAA).

[17] Apache Spark may be configured without MDSS, but doing so limits the data sources to those accessible natively via Java Database Connectivity (JDBC) and Open Database Connectivity (ODBC) enabled environments, such as IBM DB2 and IBM IMS. MDSS provides a significant set of advantages to both the application and the runtime.

programmers during installation of the *Data Service* server. *Security Optimization Management (SOM)* manages and optimizes mainframe security authentication for any Spark process that requires authentication, such as a web service or SQL call.

- ***Secure Sockets Layer (SSL)*** for the Data Service servers is transparently supported by the Application Transparent Transport Layer Security (AT-TLS) facility.
- ***Enterprise auditing*** supports the unique security requirements of Internet applications while operating in traditional enterprise computing environments.
- ***Web applications*** that access IBM z/OS operating system data and transaction can be used by analysts who do not have or do not need mainframe users IDs.
- ***Resource protection*** is provided by the *IBM Resource Access Control Facility (IBM RACF)* classes and similar ISV products.

In summary, all Spark memory structures that contain sensitive data are governed by z/OS security capabilities. Only authorized users are able to read from the z/OS data sources.

z Systems Architecture Drives Unique Apache Spark Capabilities

Enterprises now have the capability of an open-source architecture for advanced data analytics to enjoy the same qualities-of-service common to all IBM mainframe systems. Spark's in-memory-only data structures can exploit large pages, advanced memory management, data compression *(zEDC)* (especially when compressing internal data for caching and shuffling), and *Remote Direct Memory Access (RDMA)* communications to provide a high-performance scale-up and scale-out architecture. Performance is enhanced by Simultaneous Multithreading (SMT) and Single-Instruction, Multi-Data (SIMD) capabilities. Performance and costs benefits also are achieved by leveraging the *IBM System z Integrated Information Processor (zIIP)*[18].

Development and Deployment is Simplified

Analytics code developed on any platform can be used on z Systems; applications developed on other platforms can be used on z Systems, but now they have access to data that they didn't have access to before. Users can now access data on z/OS using all of the languages they already know and use, such as Java, Python, R, and Scala. There is no need to learn new skills, thus improving the productivity of developers and data scientists. And because Spark is an open source project of the Apache Foundation and not a "product", there is promise of more to come.

Conclusion

Apache Spark for z/OS delivers a more flexible, extensible, interoperable, and integrated environment for high performance, complex data analytics while avoiding the complexity, cost, and risk penalties associated with distributed ETL outboard processing environments. User-friendly should also be added to this list of benefits. But if your enterprise already is committed to a hybrid cloud model, this presents no obstacles to an Apache Spark on mainframe implementation. With Spark running on z Systems, you have the many advantages of more direct access to mainframe data via z/OS, Linux on z Systems, and the very high performance of the DB2 Accelerator (IDAA). Spark on distributed platforms – be they UNIX, Linux, or Windows-based – can still leverage z Systems data by access through JDBC (Java Database Connectivity).

The promise of Apache Spark for z/OS is to enable the exploration of all your up-to-date (i.e., most current) enterprise data – transactional, active repository, archival, structured or unstructured, whether or not IBM branded –all within one or two frames. It features an open source integrated architecture that is focused on high performance and functionality, scalability, security, and the resiliency of IBM mainframe technology. Moreover, it fits easily into the existing data analytics infrastructure, both hardware and software, already established by IBM. The challenge now is yours – you need do a thorough investigation of Apache Spark for z/OS as your enterprise considers how it most productively can widen the focus of your enterprise's missions and deliver current data and up-to-date answers when they are needed.

[18] The zIIP (z Systems Integrated Information Processor) is a special purpose processor that accelerates a number of functions including Java Virtual Machine (JVM) task and Java application code, z/OS SML System Services, several DB2 for z/OS functions, among others. When engaged, in most cases transparently, is lowers the cost of computing by off-loading several z/OS processes.

About The Clipper Group, Inc.

The Clipper Group, Inc., now in its twenty-fourth year, is an independent publishing and consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

Stephen D. (Steve) Bartlett **is a Senior Contributing Analyst for The Clipper Group.** Mr. Bartlett's interests include enterprise solutions including servers, system software, middleware, their underlying technologies, and their optimal deployment for responsiveness to emerging business requirements. In 2010, he joined the Clipper team after over 42 years with the IBM Corporation as an account and program manager in large system sales, product development, strategy, marketing, market research, and finance. During that time, he received several awards for his contributions in innovative market research and contributions to the business. He has a B.S. from Rensselaer Polytechnic Institute, and an M.S. from Union College.

- ***Reach Steve Bartlett via e-mail at steve.bartlett@clipper.com or at 845-452-4111.***

Regarding Trademarks and Service Marks

Disclosures

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

After publication of a bulletin on *clipper.com*, The Clipper Group offers all vendors and users the opportunity to license its publications for a fee, since linking to Clipper's web pages, posting of Clipper documents on other's websites, and printing of hard-copy reprints is not allowed without payment of related fee(s). Less than half of our publications are licensed in this way. In addition, analysts regularly receive briefings from many vendors. Occasionally, Clipper analysts' travel and/or lodging expenses and/or conference fees have been subsidized by a vendor, in order to participate in briefings. The Clipper Group does not charge any professional fees to participate in these information-gathering events. In addition, some vendors sometime provide binders, USB drives containing presentations, and other conference-related paraphernalia to Clipper's analysts.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.