

# Hadoop and System z



An IBM Redbooks Point-of-View publication from the IBM Competitive Project Office

By **Vic S. Leith**, Sr. Software Engineer, and **John J. Thomas**, Sr. Technical Staff Member

## Highlights

- ▶ Thanks to their power and flexibility, a wide range of use cases involving Apache Hadoop and IBM System z are emerging today.
- ▶ Integrating System z data with on-premises and off-premises Hadoop environments presents challenges, particularly in the areas of data governance and data ingestion.
- ▶ Recent IBM testing has proven that System z governance can be extended to hybrid Hadoop implementations, with System z policies in control from end to end.
- ▶ Multiple hybrid models were tested, so clients can choose the one that best addresses their business and IT needs.

## Emerging use cases

A large portion of the world's enterprise data is stored on IBM® System z® mainframes, and enterprises have long relied on this structured data for their day-to-day operations. At the same time, exponential growth in mobile technologies, social media, and other systems of engagement is driving the emergence of new forms of data, both unstructured and semi-structured. Combining operational data with these newer data sources can yield deep business insights. For example, companies can combine their operational data with collected social media sentiments to create a 360-degree view of their customers, which can provide a key competitive edge when identifying new opportunities.

Apache Hadoop is an open source software framework that has become a leading method of processing various data types, particularly semi-structured and unstructured data. And, an IBM implementation of Hadoop, which is called IBM InfoSphere® BigInsights™, is available on multiple platforms.

So, how can System z and z/OS®-based enterprises take advantage of the power of Hadoop? There is, in fact, a wide spectrum of use cases linking Hadoop processing with System z. This IBM Redbooks® Point-of-View publication focuses on two use case categories where System z-Hadoop integration has proven successful.

The first category deals with structured and semi-structured data that originates on System z. Operational data from z/OS in the form of logs, XML or VSAM files, DB2® data, IMS™ data, and so on can be brought into a Hadoop environment and processed by using Hadoop tools. The result set can be consumed by System z or used elsewhere downstream.

Figure 1 is an example of this use case, showing a scenario where System z data is processed by using Hadoop.

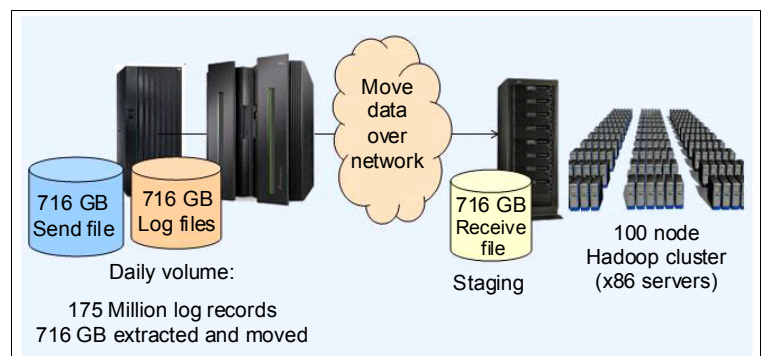


Figure 1 Using Hadoop to process data originating on System z



The second category, which can exist alone or in combination with the first category, deals with large amounts of semi-structured or unstructured data that originates outside System z but must be integrated with operational data on System z. Again, the processing can be done in the Hadoop environment and the result set can be consumed by System z or downstream.

Figure 2 provides an example of this second use case category, showing a scenario in which external data in Hadoop is integrated with System z operational data.

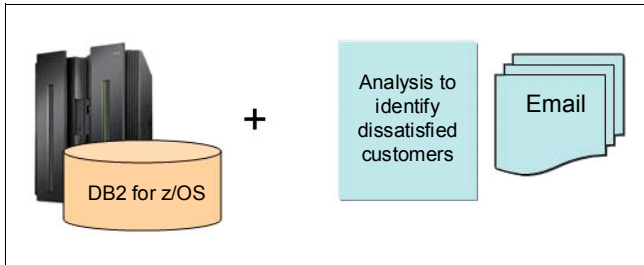


Figure 2 Integrating external Hadoop data with System z operational data

## Extending System z governance and data ingestion to Hadoop

Distributed Hadoop implementations are widespread. However, clients sometimes face challenges creating and using a hybrid environment that includes System z, z/OS, and Hadoop. Two of the most common challenges are described:

- ▶ **Data governance:** System z administrators have well-established standards that govern processing, security, and regulatory and auditing requirements, but in each category, System z needs to be in control of the data. So, how do System z administrators address these operational practices in a Hadoop context? How can existing System z security policies be extended to a Hadoop environment?
- ▶ **Ingestion of data:** To move data from System z into the Hadoop environment, you need high speed, optimized connectors between traditional z/OS logical partitions (LPARs) and a System z-controlled Hadoop environment. How and when is System z data moved into Hadoop?

There are several configurations that address the challenges of Hadoop processing in conjunction with System z:

- ▶ **On-premises distributed hybrid:** An on-premises Hadoop cluster, built on discrete hardware, under System z control
- ▶ **On-premises Linux on System z hybrid:** An on-premises Hadoop cluster built with Linux on System z nodes
- ▶ **Off-premises distributed hybrid:** An off-premises Hadoop cluster built with appliances, discrete nodes, or cloud-based nodes, under System z control

In each of these examples, two key requirements exist:

- ▶ The ability to efficiently extract and load data
- ▶ End-to-end processing capability, including data loading, job submission, and results retrieval using System z credentials and under System z governance

This paper describes approaches to successfully construct and use on-premises or off-premises Hadoop environments and govern those environments using end-to-end System z policies. Considerations include governance when loading data into Hadoop, controlling who can load that data, limiting the data that can be loaded, ensuring adequate security around submitting map reduce jobs, monitoring job progress from System z, and retrieving the results.

## Three proven hybrid Hadoop options

IBM conducted a series of tests that studied several System z-governed hybrid Hadoop options. These environments are depicted in Figure 3 on page 3 and listed here:

- ▶ **On-premises distributed hybrid**  
This environment consists of a z/OS LPAR with an on-premises distributed Hadoop cluster connected directly to System z. The Hadoop cluster can be dedicated hardware or a prebuilt Hadoop appliance.
- ▶ **On-premises Linux on System z hybrid**  
This environment consists of a z/OS LPAR and a multi-node Hadoop cluster running as Linux on System z guests. To provide for this option, IBM recently announced IBM InfoSphere BigInsights for Linux on System z.

► **Off-premises distributed hybrid**

This environment consists of a z/OS LPAR with an off-premises distributed Hadoop cluster connected using a secure VPN tunnel. These clusters can be appliances, discrete nodes, or cloud-based nodes.

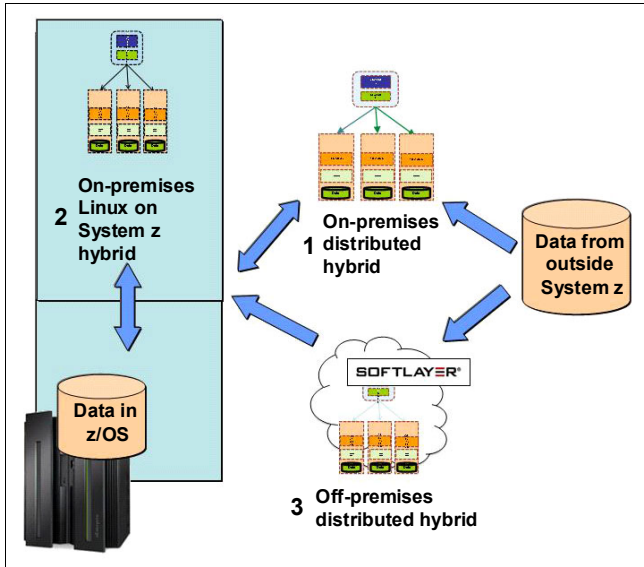


Figure 3 System z-governed Hadoop environment options

A System z-governed distributed hybrid environment requires System z-controlled, secure communications between the mainframe and the distributed Hadoop clusters. Test environment 2 used the IBM InfoSphere System z Connector for Hadoop to communicate between z/OS and the Hadoop cluster. Test environments 1 and 3 used the Co:Z Co-Processing Toolkit from Dovetail Technologies for this communication. The Co:Z toolkit consists of a streaming library on the z/OS side and a number of distributed platform agents. The IBM InfoSphere System z Connector for Hadoop provides a graphical interface for moving z/OS data into the Hadoop environment on Linux for System z. The Co:Z library fits into the z/OS batch processing environment and allows the use of simple batch jobs to accomplish the tasks.

**On-premises distributed hybrid environment**

Figure 4 shows the first on-premises distributed hybrid environment that we tested. It consisted of a z/OS LPAR and a five-node, local x86-based InfoSphere BigInsights Hadoop cluster.

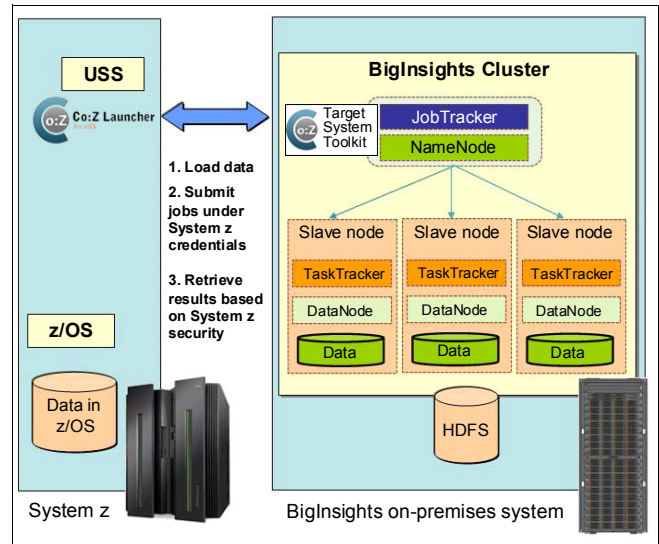


Figure 4 On-premises distributed hybrid environment

Two types of on-premises tests were executed:

- End-to-end: This test compared the complete set of steps needed to accomplish an end-to-end data analytics job.
- Data load: This text compared the steps needed to transfer a file from System z into the Hadoop Distributed File System (HDFS).

**End-to-end test**

The end-to-end test consisted of five steps:

1. Transfer a VSAM file from System z to a local distributed Hadoop master node's Linux ext4 file system (because the file transfer must be secure, FTP is not recommended as a transfer option).
2. Write the file into HDFS by using the command `hadoop fs -put`.
3. Use BigInsights to execute a Hadoop map reduce job to analyze the data.
4. Copy the result set from HDFS to the Linux file system by using the command `hadoop fs -get`.
5. Return the result set to System z.

Figure 5 on page 4 shows the end-to-end on-premises usage scenario.

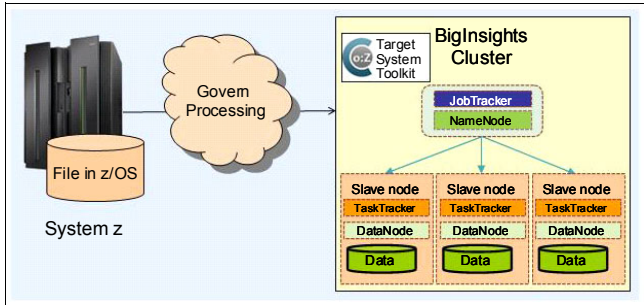


Figure 5 On-premises usage scenario: Hadoop analysis of existing VSAM file or DB2 extract

In this end-to-end test, a single Job Control Language (JCL) batch job performed all of the necessary processing to transfer the VSAM file from System z to HDFS, execute the map reduce job, and return the result set to System z.

The batch job included three Co:Z steps:

1. It moved the VSAM file from z/OS across the network and loaded it using Co:Z pipes into the HDFS. Unlike a file transfer using Secure Shell (SSH) or Secure FTP (SFTP), Co:Z processing allowed the network transfer and the write into the Hadoop file system to occur without creating a temporary file on the Linux ext4 file system.
2. It issued the command to perform the map reduce job.
3. It moved the result from HDFS back into z/OS, again, without requiring a temporary file on the Linux ext4 file system to get it out of Hadoop.

For this end-to-end scenario, Co:Z only requires the SSH keys that are pre-loaded into the Hadoop master node. All of the processing is directed by z/OS; therefore, a Resource Access Control Facility (RACF®) PassTicket is not needed. z/OS uses Co:Z features to execute remote jobs on the Hadoop master node.

### Data load test

The data load test consisted of two steps:

1. Transfer data from System z to a local distributed Hadoop master node's Linux file system (which serves to test real-time data extraction directly from DB2 tables).
2. Write the data into HDFS using the Hadoop command `hadoop fs -put`.

Figure 6 shows the data load on-premises usage scenario.

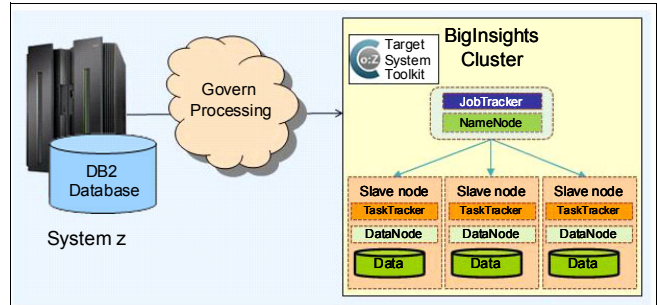


Figure 6 On-premises usage scenario: Processing System z DB2 data directly

In this data load test, a single JCL batch job performed all of the processing necessary to extract data from DB2 using High Performance Unload (HPU) and then to transfer that data from System z to HDFS. The job uses HPU to extract data, transfer it, and load it into HDFS. The data is never written into a System z file or a Linux file outside of HDFS. Co:Z extracts the data, transfers it across the network, and loads it directly into the HDFS file system by using Co:Z pipes.

In addition, to demonstrate the ability to govern end-to-end processing using System z security policy models, we executed a System z action that was initiated from the remote Hadoop master node over a Co:Z connection. In this scenario, System z issues a command to the remote Hadoop master node that requires remote initiation of a System z job to complete its processing. A RACF PassTicket is required to allow the remote access to System z resources.

### On-premises Linux on System z hybrid environment

An on-premises Linux on System z hybrid environment is ideal for Hadoop use cases where dataset sizes are relatively small and the data being processed originates entirely on System z.

Figure 7 on page 5 shows the environment that we tested, which consisted of a z/OS LPAR and a five-node local BigInsights for Linux on System z Hadoop cluster implemented by using Linux on System z virtual machines. Communications with z/OS were handled using the IBM InfoSphere System z Connector for Hadoop that is included in the BigInsights for Linux on System z offering. This connector easily transfers data from z/OS sources while still ensuring security when running Hadoop on the mainframe. It provides a graphical interface for data transfer that minimizes any need for system programming or specialized skills.



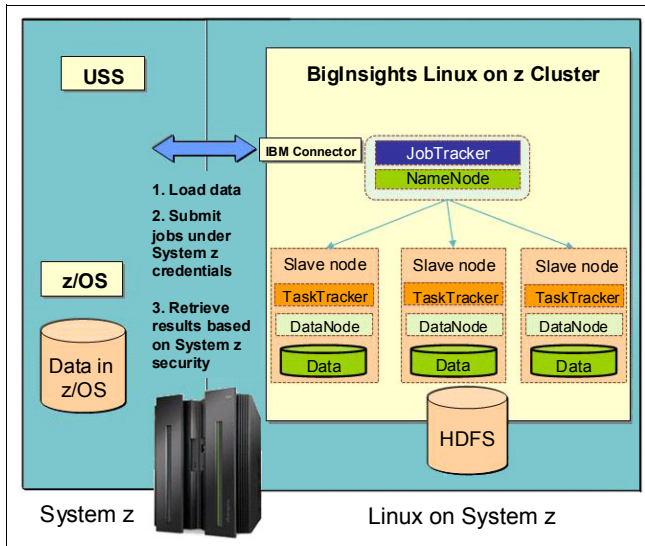


Figure 7 On-premises Linux on System z hybrid environment

The same end-to-end and data load tests that were performed on the on-premises distributed environment were performed on this on-premises Linux on System z hybrid environment, with equally successful results.

### Off-premises distributed hybrid environment

An off-premises distributed hybrid model is ideal if you need to process bulk data from external sources without bringing it all in-house. The off-premises distributed hybrid environment we tested consisted of a z/OS LPAR and a five-node, x86-based BigInsights Hadoop cluster implemented using IBM SoftLayer, the IBM public cloud offering. Figure 8 shows a high-level view of this type of off-premises environment.

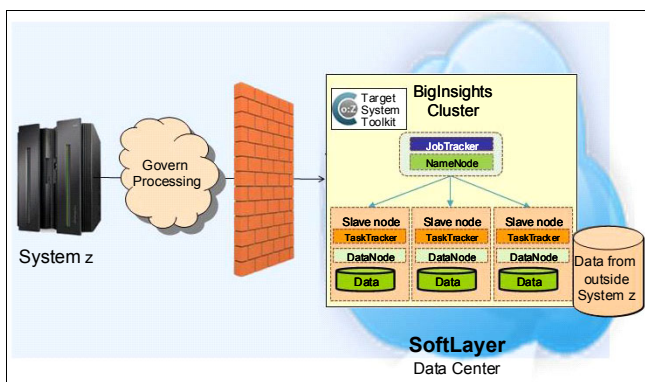


Figure 8 Off-premises distributed hybrid environment

One important aspect of the SoftLayer integration is the use of a private VPN between on-premises servers and the SoftLayer data center. By using the SoftLayer portal, requesting and deploying servers on SoftLayer is relatively easy. After the SoftLayer servers are available, creating a SoftLayer BigInsights Hadoop cluster on those servers is simple.

The SoftLayer Hadoop nodes' public IP addresses were only used during the test's server installation steps. For security reasons, after the Hadoop servers were operational, their public IP addresses were disabled. The SoftLayer Hadoop servers were then only accessible through the private VPN.

This off-premises configuration was subjected to a data analysis test. The test included the commands used to accomplish a sample data analytics job. The data analysis test consisted of three steps:

1. Execute a map reduce job to analyze the data in an off-premises BigInsights cluster.
2. Copy the result set from HDFS into the Linux file system by using the Hadoop command `hadoop fs -get`.
3. Return the result set to System z.

Again, a single JCL batch job performed all of the necessary processing to execute the map reduce job and to return the result set to System z.

The batch job included two Co:Z steps:

1. It issued the command to perform the map reduce job on the off-premises Hadoop cluster.
2. It moved the result from HDFS back into z/OS (in other words, out of Hadoop) without requiring a temporary file on the Linux ext4 file system.

For this data analytics scenario, Co:Z only required SSH keys that are pre-loaded into the Hadoop master node. All of the processing is directed by z/OS; no RACF PassTicket is necessary. z/OS uses Co:Z features to execute remote jobs on the Hadoop master node.

## Conclusions

Our tests proved several points about integrating System z and Hadoop and illustrated several benefits:

- ▶ Enterprises can construct and use the on-premises or off-premises Hadoop environments that best fit their needs, and enterprises can govern those environments by using end-to-end System z policies.
- ▶ In these environments, enterprises can use Hadoop to make greater use of their traditional System z data while also taking advantage of newer, emerging forms of data to derive greater business value.

## What's next: How IBM can help

IBM makes it simple to use Hadoop to get new business value by combining operational data with newer data sources.

IBM InfoSphere BigInsights for Hadoop empowers enterprises of all sizes to cost-effectively manage and analyze the massive volume, variety, and velocity of new forms of data that consumers and businesses create every day.

By using the techniques described in this paper, you can create System z-controlled hybrid Hadoop environments that best fit the needs of your workload.

## Resources for more information

For more information about the concepts highlighted in this document, see the following resources:

- ▶ IBM InfoSphere BigInsights:  
<http://www.ibm.com/software/data/infosphere/biginsights>
- ▶ IBM InfoSphere BigInsights for Linux on System z and IBM InfoSphere System z Connector for Hadoop:  
<http://www-01.ibm.com/software/os/systemz/biginsightsz/>
- ▶ Apache Hadoop:  
<http://hadoop.apache.org/>
- ▶ Dovetail Technologies Co:Z Co-Processing Toolkit:  
<http://dovetail.com/solutions.html#coz>
- ▶ IBM SoftLayer  
<http://www.softlayer.com/>

# Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to: *IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.*

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:** INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

## COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

This document, REDP-5142-00, was created or updated on October 15, 2014.



## Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

BigInsights™  
DB2®  
IBM®  
IMS™  
InfoSphere®  
RACF®  
Redbooks®  
Redbooks (logo)   
System z®  
z/OS®

The following terms are trademarks of other companies:

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.