

The Impact of Hadoop Resiliency on SAS® LASR™ Analytic Server

Rob Collum, SAS Institute Inc.

ABSTRACT

The SAS® LASR™ Analytic Server acts as a back-end, in-memory analytics engine for solutions such as SAS® Visual Analytics and SAS® Visual Statistics. It is designed to exist in a massively scalable, distributed environment, often alongside Hadoop. This paper guides you through the impacts of the architecture decisions shared by both software applications and what they specifically mean for SAS®. You will also see positive actions you can take to rebound from unexpected outages and resume efficient operations.

INTRODUCTION

What happens if you lose a data node from your Hadoop cluster? How do you recover? What impact will this have on the performance of SAS Visual Analytics?

Like a lot of things, the answers will depend on the specifics of your site and solution implementation. So let's look at what goes on when a Hadoop Data Node goes down in an environment where the SAS LASR Analytic Server is also deployed.

WHAT'S WHAT

The SAS Visual Analytics (VA) solution bundle for distributed computing environments includes a couple of key pieces of software that are relevant to the topic of this paper:

1. SAS LASR Analytic Server: the in-memory engine that provides extremely fast and scalable analytics performance. It's the number-crunching brain behind SAS Visual Analytics.
2. SAS High-Performance Analytics Deployment of Hadoop (HPDH): built from the same package of Apache Hadoop that is freely downloadable from the Apache website. The only thing "extra" is that SAS includes a couple of additional JAR files that enable it to provide the SASHDAT capability. Of course, those JAR files are also available for use with other supported distributions of Hadoop, but it's convenient in this paper to assume that HPDH is in play.

It's pretty common for a deployment of VA to also include a co-located installation of HPDH. This allows the SAS LASR Analytic Server—which is the in-memory analytics engine—to take advantage of the speed and efficiency of using SASHDAT tables when lifting data from disk up to RAM.

Figure 1 illustrates LASR with Hadoop as a co-located data provider.

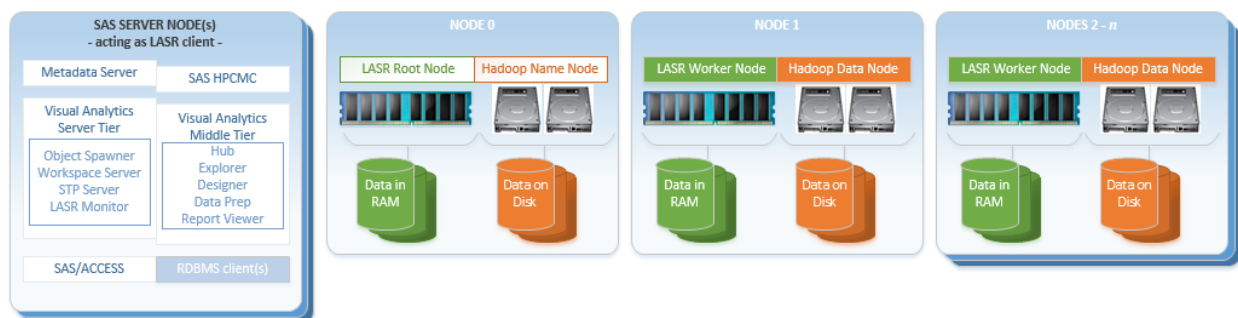


Figure 1. Co-located Architecture

This is a good time to point out that the Hadoop technology is primarily focused on disk- (or SSD-) based storage. As a software system, Hadoop will, of course, consume RAM, CPU, and network resources, but those are not significant when compared to the reliance of Hadoop on physical disk space.

Further, the SAS LASR Analytic Server is primarily focused on RAM-based storage. Like other software systems, it will consume disk, CPU, and network as well, but RAM is the emphasis and constraining factor.

Due to their respective hardware priorities, co-locating Hadoop and LASR is a complementary arrangement that is often beneficial in terms of hardware utilization, efficiency, and cost.

This paper will focus on a co-located deployment of LASR and Hadoop. It's worth noting that a co-located deployment is not the only possible topology – separately hosted clusters for Hadoop and LASR are also supported. However, this co-located architecture is commonplace and requires specific considerations discussed here.

NAME NODE AND ROOT NODE

The Hadoop Name Node is in charge of the Hadoop software services in the cluster—specifically for this paper, the Hadoop Distributed File System (HDFS). The Name Node keeps track of all activities going on in HDFS and manages the cluster to ensure that data availability is maintained.

The LASR Root Node is in charge of the SAS LASR Analytic Server services. It handles incoming requests, directs the activity of all in-memory analytics processing, and returns the results.

In the figure above, the Hadoop Name Node and LASR Root Node are placed on the same host machine. This is convenient for our discussion, but is required when utilizing Hadoop as a co-located data provider for LASR.

DATA NODES AND WORKER NODES

The Hadoop Data Nodes are each independent services that reside on their own host machines. They do not share any hardware resources with their peers. The Data Nodes are responsible for storing data on disk, so the disk storage system is the critical hardware component to follow.

The LASR Worker Nodes are also each independent services that reside on their own host machines. Nor do they share any hardware resources with their peers. The Worker Nodes are responsible for performing analytic functions on data which has been lifted into each machines' memory. Therefore, RAM is the critical hardware component to follow.

In the figure above, each Hadoop Data Node and LASR Worker Node share a common host machine. This is convenient as well as required when utilizing Hadoop as a co-located data provider for LASR.

BASIC HADOOP PROCESS

When a file must be saved to HDFS, it's broken up into chunks called "blocks." The Hadoop Name Node assigns blocks to the various Data Nodes in the cluster. To ensure availability of the data in the face of a hardware failure, each block is replicated, usually three times. Therefore, the replication factor has a direct impact on the amount of disk space used. To store 1 terabyte of data in HDFS with a replication factor of 3, there must be a total of 3 terabytes of disk space available in the cluster. Of course, the replication factor is configurable.

Further, each block is placed on a different Data Node than its replicated sibling blocks. That way, if one Data Node fails for some reason, there are two others that can satisfy the request for those blocks. Hadoop can even be configured for rack-awareness. This means that in a multi-rack environment,

Hadoop will distribute the data blocks to servers on multiple racks to protect against the loss of an entire rack of servers.

SASHDAT

The SASHDAT data format is a binary, compressible storage structure. It is optimized for analytics processing in high-performance environments with concurrency—like Hadoop!

When SASHDAT tables are created in HDFS, the blocks are evenly distributed across all of the LASR Worker node machines to ensure that each node hosts the same amount of data.

By default, SAS directs Hadoop to use a replication factor of only 2 for SASHDAT. This option is also configurable.

LOADING LASR WITH DATA

A LASR Root Node and a group of LASR Workers distributed across multiple machines form a logical SAS LASR Analytic Server.

The SAS LASR Analytic Server is a persistent, in-memory analytic engine. This means it's designed to process large volumes of data that have been completely lifted up into RAM. After the data is in RAM, there's no waiting for the disk anymore. This makes LASR extremely fast and efficient.

Multiple tables can be added to a SAS LASR Analytic Server—as many as RAM will allow (with allowance for other RAM-dependent processing). Also, multiple SAS Analytic LASR Servers can be started. Each one acts independently of the others.

The data loaded in LASR can literally come from anywhere that Base SAS® (and its suite of SAS/ACCESS® engines) can access it. Loading data into LASR in those scenarios is typically a serial process where the data flows to the LASR Root Node, which then distributes it equally to the LASR Workers.

Figure 2. "Serially Loading Source Data into the SAS LASR Analytic Server," shows this concept in action.

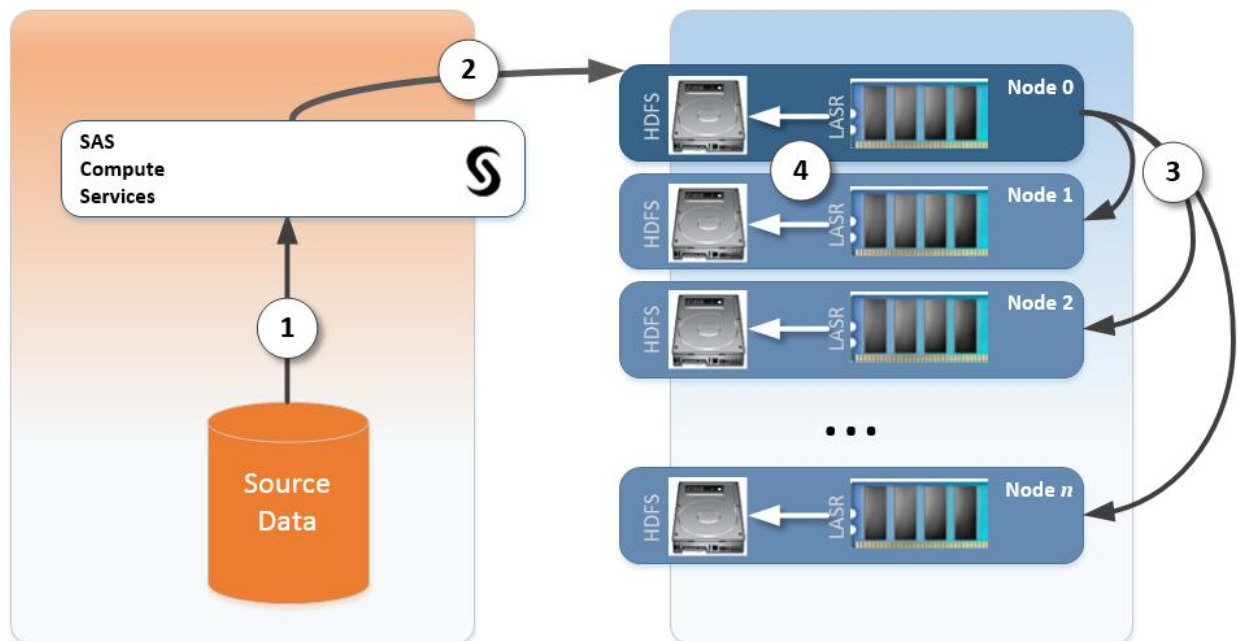


Figure 2. Serially Loading Source Data into the SAS LASR Analytic Server

The steps illustrated include the following:

1. SAS software reads the data from its source (SAS data sets, raw text files, or third-party data via SAS/ACCESS).
2. SAS sends the data over to the Root Node of the SAS LASR Analytic Server.
3. The LASR Root Node then distributes the data evenly to each of its Workers.

This approach works great, but the serial-loading approach can create a bottleneck at the Root Node when you are working with really large data. We can therefore take one more step:

4. Push the data from each LASR Worker down to the co-located Hadoop Data Node in HDFS as SASHDAT tables.

Having this data staged in SASHDAT format then provides the fastest and most efficient way to load very large volumes of data out of Hadoop and back into LASR. Leveraging SASHDAT, the data can be lifted from disk by each Hadoop Data Node directly into RAM to each co-located LASR Worker Node. This technique provides multiple parallel streams. The result is very fast LASR load times as compared to the serial technique. While LASR is a persistent service, it is *not* permanent. And while our MPP environment provides a lot of RAM, it's not infinite. Having this ability to drop tables from LASR and then reload them quickly on demand is very important.

Figure 3 shows the parallel lift of SASHDAT data into RAM for the LASR Workers.

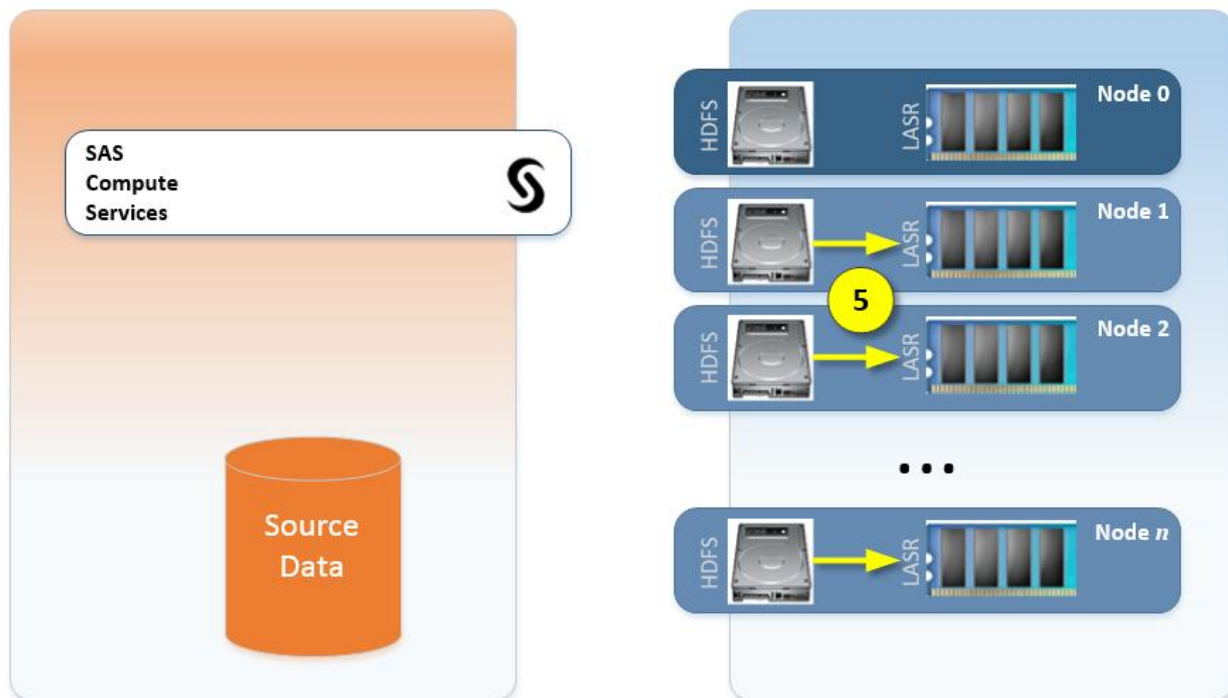


Figure 3. Parallel Loading SASHDAT Data into the SAS LASR Analytic Server

The figure illustrates the following action:

- The data stored in SASHDAT is evenly distributed across all nodes of the co-located data provider (HDFS). Upon request, those blocks are lifted directly into RAM from each Hadoop Data Node to the associated LASR Worker.

LASR PROCESSING THE DATA

When dealing with requests from the outside world, the LASR Root Node handles most communication with the outside software services and coordinates the actions of the LASR Workers. When a request comes in, the LASR Root Node directs all of the workers to perform analytics on their respective chunks of data. The interim results from each Worker are then sent up to the Root for final summation and processing. The Root then returns the final result set in response to the initial request.

Figure 4 shows how LASR deals with incoming requests.

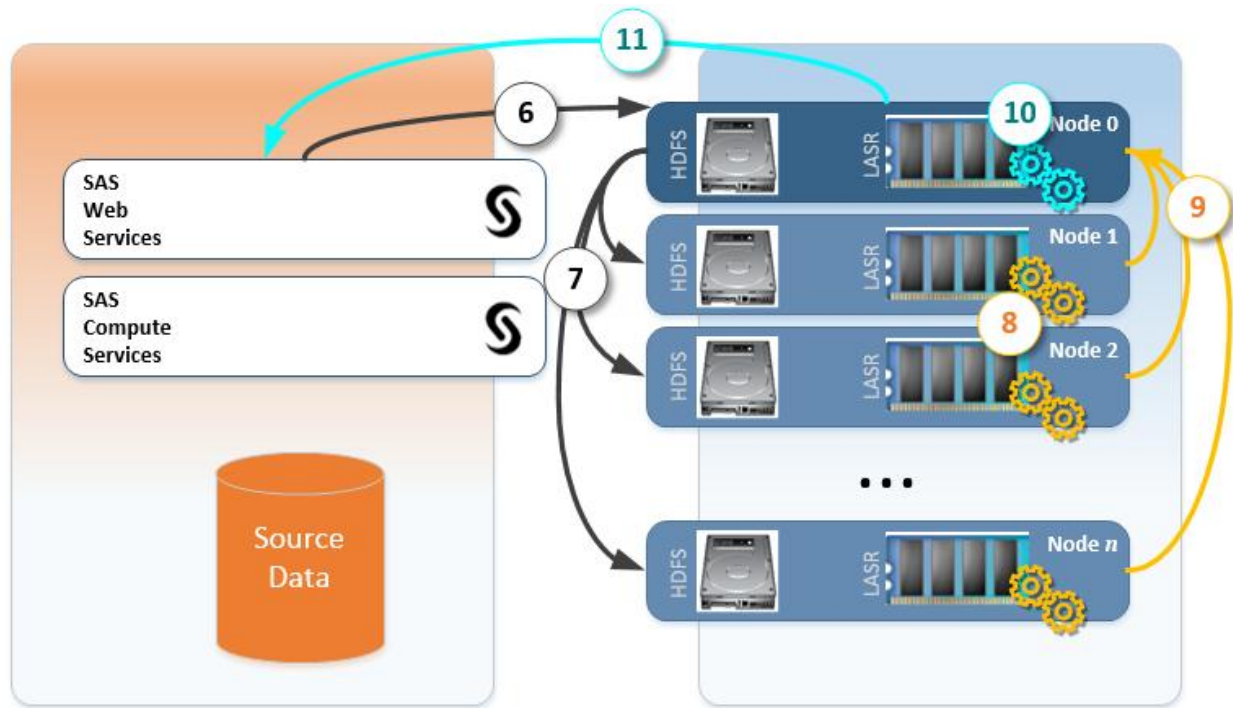


Figure 4. Requesting Analytics from LASR

The steps illustrated in the diagram show the following:

- A user (not shown above) wants to view a report. A SAS web application (like SAS Visual Analytics—but there are other sources) requests LASR to perform analysis on a table currently in memory so that it can render the result.
- The LASR Root Node directs the LASR Workers to perform the job.
- The LASR Workers run the necessary calculations on the portion of data contained on that node.
- Each LASR Worker returns the results of its calculations back to the LASR Root.
- The LASR Root Node waits for all of the Workers to return their results and then calculates the final answer.
- The LASR Root replies with the result dataset to the requestor, and the user (not shown in the diagram) sees the report.

LASR ASSUMPTIONS

We've established several important architectural criteria about the cluster environment, the data distribution, and the workload up to this point.

LASR is designed to assume that each Worker Node is built with similar hardware specifications (same CPU, same RAM, and so on). Further, data is intentionally distributed to each of the Workers in equal portions. Therefore, each Worker Node should be able to complete a given request to process the same amount of data in approximately the same amount of time as all of the other Worker Nodes. On the other hand, for example, if just one Worker Node must process more data than any of the others, then the entire result is delayed. With this in mind, we often say that the SAS LASR Analytic Server is only as fast as the slowest node.

WHEN A HOST MACHINE IN THE CLUSTER FAILS

So now's the time to cover a possible failure mode that we're likely to see in a distributed LASR Analytic Server with Hadoop as a co-located data provider. That is, Hadoop and LASR share a common hosting environment which acts as the analytics processing backend of our SAS Visual Analytics solution deployment.

We'll start here with the impact on Hadoop and then expand later into the ramifications for LASR.

ONE DATA NODE DOWN

So what happens to Hadoop if we lose a Data Node? Maybe the HDFS service was killed on a particular machine. Or maybe a hard drive failed, and it lost its HDFS blocks. For this example, suppose that the power supply got fried in that rack server and that one machine has gone dark—no OS or other software services are running on the failed host.

As far as Hadoop is concerned, there's no really critical problem. Even with the complete loss of a single Data Node, you can request that HDFS provide you with your file, and it will come back just fine. That's because Hadoop can find the replicated blocks on other Data Nodes that are copies of the primary blocks that were stored on the lost Data Node. To an end user (or service), there's no apparent failure: you asked for your file from HDFS and got it back.

Hadoop is specifically designed to *expect* hardware failures. This design allows the software to work around the lost Data Node, providing you with a highly available data storage service.

HOW HADOOP DEALS WITH IT

In order for Hadoop to maintain the façade of uninterrupted availability, it must perform some actions behind the scenes.

In this example, the Data Node remains offline for more than 10 minutes. The Hadoop default behavior then is to remove that machine from its list of active servers. After that, Hadoop won't even try to work with the failed machine anymore.

Remember the replicated blocks? With that Data Node offline, Hadoop now no longer has the minimum number of replicated blocks to ensure continued availability. Because the Name Node keeps a detailed audit of exactly which blocks are where, it is aware of the shortage. So it will then direct the remaining Data Nodes to make additional replicated blocks among themselves. That way, the minimum expectations for performance and availability can be re-established with the rest of the cluster.

If no further action is taken, Hadoop will happily continue to function minus the one Data Node, assuming that it has enough hardware remaining in the cluster to continue full availability operations.

BUT WHAT ABOUT LASR?

Our LASR Workers are co-located with the Hadoop Data Nodes. So when that one host machine died, we not only lost the Hadoop services on that machine, we also lost any running LASR Worker services.

A logical SAS LASR Analytic Server is not robust enough to seamlessly work around the loss of a LASR Worker node. In our failure scenario, the loss of one physical host machine has corrupted the running logical SAS LASR Analytic Server—it lost the subset of data that it had in RAM. There is no option but to manually stop the entire SAS LASR Analytic Server.

With LASR not running, the data it was responsible for isn't available. This means that SAS Visual Analytics cannot serve up any reports that rely on that data.

Fortunately, we have the option of re-starting the LASR process. As the Root and remaining Worker Nodes come online, they will automatically form up into the same logical SAS LASR Analytic Server as before, with the difference being that now one node is missing.

After our SAS LASR Analytic Server is back online, we can then reload data into it. All of the data we need is still available as SASHDAT tables in HDFS thanks to the built-in availability design of Hadoop.

SITUATION REPORT

At this point of our failure scenario, it's a mixed bag of good and bad news:

1. The physical rack server machine has a faulty power supply and is off.
2. Hadoop has detected the loss of its Data Node hosted on that server. It has removed the Data Node from its list of Active Servers and created additional replicated data blocks, which are distributed to the other Data Nodes to ensure continued availability in the face of another outage.
3. The logical SAS LASR Analytic Server (Root Node + Worker Nodes) was corrupted due to the loss of one Worker. There was a brief service outage when we manually stopped and restarted the SAS LASR Analytic Server and then reloaded data into it from SASHDAT.
4. With Hadoop re-replicating data blocks for HDFS, the original, even distribution of data we had when the SASHDAT tables were created is no longer possible. Hadoop follows its own set of criteria for determining block distribution, putting more weight toward uniform disk utilization. This means that some LASR Workers will have more data to process than others, further exacerbating the decreased performance caused by the loss of a physical machine.
5. While performance is certainly degraded, the final word is that, with only a brief interruption to LASR service, we can continue to provide back-end data and in-memory analytic services to SAS Visual Analytics.

GETTING BACK WHOLE

Of course, we're not going to leave the failed blade server just lying there dark on the rack. One of our industrious IT staff will pull the unit and replace the power supply, which arrived in priority mail over the weekend. After some testing, the server machine gets slotted back into the rack and switched on to resume work.

HADOOP

Because the failed server was automatically removed from the list of active Hadoop servers, we must inform Hadoop that the freshly restored Data Node is back and available for work. Hadoop makes this painless and can perform the task without any interruption in service.

1. Ensure the node's DNS name is properly referenced in the following:

- The `conf/slaves` file on the Name Node machine (for the start-up scripts on the Name Node)
 - The `dfs.include` (or `.exclude`) file on the Name Node machine (notify HDFS directly of active nodes)
 - If changed, then execute `$HADOOP_HOME/bin/hadoop dfsadmin -refreshNodes`
2. Log on to the newly restored machine with the Hadoop administrator account.
 3. Start the Data Node: `$HADOOP_HOME/bin/hadoop-daemon.sh start datanode.`
 4. Start the Task Tracker: `$HADOOP_HOME/bin/hadoop-daemon.sh start tasktracker.`

The Data Node will immediately be available to take on workload. If new tables are loaded, or if replication activity is underway, or some other distributed Hadoop work, then the Data Node can participate as directed.

LASR

The addition of a physical machine to the environment will not have an immediate effect on any existing SAS LASR Analytic Servers. To add this machine's horsepower to LASR, you will need to again stop and restart the SAS LASR Analytic Server so that it picks up a LASR Worker on the newly restored machine. You will then need to reload data again.

Of course, this can take place at a convenient time when a brief LASR outage will have no impact on the SAS Visual Analytics user community.

OPTIMIZING THE RESTORED SYSTEM

We've seen that Hadoop and LASR can accommodate the loss of a physical machine in the cluster. Further, we've seen that re-introducing the node into the cluster is pretty much painless. The great thing is that the Hadoop and LASR services can operate in these different conditions with little overt impact on the end users.

But there's more that we can do, especially for LASR. The distribution of blocks for the SASHDAT tables in HDFS has been changed by the Hadoop availability procedures. This means that some Data Nodes will now have a greater number of blocks than others. And then, when that data is lifted up into RAM to the associated LASR Workers, some of those Workers will have more data to process than others. This will cause slightly longer analytic processing times. Fortunately, there is something you can do to help with this situation.

DISK UTILIZATION

One of the many cluster attributes that Hadoop monitors and adjusts for is disk utilization of the nodes in the cluster. When Hadoop distributes blocks in HDFS, it ensures that those blocks are equitably rationed to each node while also making adjustments for availability concerns based on rack awareness, neighboring servers, and even the specific disk. The goal is to ensure that data is available in the event of a physical interruption to a disk, a server, a rack, the network, and so on. The end result is that each node should utilize about the same amount of disk space.

To be clear, equitable disk utilization isn't the primary goal, just an important one. If a brand new node is added to the Hadoop cluster, its disk utilization will be near zero. Hadoop is designed so that it will not monopolize the new node in an attempt to "catch up" its disk utilization with the other nodes in the cluster.

With this in mind, you can see that Hadoop isn't concerned with the equal distribution of data in the environment. Its goal is more toward equal disk utilization, which is affected by many factors. Achieving one of these—utilization versus distribution—will often yield a close approximation of the other, but with varying success.

DATA DISTRIBUTION

The equal distribution of data in SASHDAT tables across the Hadoop Data Nodes is our primary concern at this point. Ensuring that each LASR Worker has exactly the same workload as its peers is an important factor in maximizing the speed and efficiency of the LASR Analytic Server.

The following table illustrates this concept. It compares how a 5-million record SASHDAT table's data might be distributed under three different conditions: (1) when the entire HDFS cluster is online, (2) after the loss of a Data Node, and (3) after recovery of the lost node.

			ALL NODES ONLINE	AFTER LOSS OF ONE NODE	LOST NODE RESTORED - ALL NODES BACK ONLINE
Machine	Hadoop Role	LASR Role	# of records per node	# of records per node	# of records per node
Node0.example.com	Name	Root	0 (meta only)	0 (meta only)	0 (meta only)
Node1.example.com	Data	Worker	~1,000,000	0	0
Node2.example.com	Data	Worker	~1,000,000	~1,400,000	~1,400,000
Node3.example.com	Data	Worker	~1,000,000	~1,200,000	~1,200,000
Node4.example.com	Data	Worker	~1,000,000	~1,100,000	~1,100,000
Node5.example.com	Data	Worker	~1,000,000	~1,300,000	~1,300,000

Table 1. Comparison of Distribution under Three Different Conditions

Notes:

- These numbers are for illustrative purposes only. Distribution of records is subject to a number of factors. It could be worse (or better) than this.
- The loss of a LASR Worker Node in this 6-machine cluster reduces analytic performance by about 20%.
- After Node1 went down, HDFS redistributed the blocks of the SASHDAT table so that it could maintain high availability. In doing so, HDFS followed its guidelines for maintaining an even utilization of disk space across the remaining nodes.
- The resulting uneven distribution of data to the remaining LASR Workers puts more workload on Node2. Node2 has more records to process, so it takes longer to complete a request.
- The LASR Root Node cannot respond to the original request until all Workers have responded. It will wait longest for Node2 to complete its portion of processing.
- After restoring Node1 to active service, Hadoop doesn't take any automatic action to redistribute the HDFS blocks again. So while the associated LASR Worker on Node1 can be started as part of a logical SAS LASR Analytic Server, it won't automatically have any data to process.

So what should we do if we want to re-establish an even distribution of SASHDAT data across all of the Hadoop Data and LASR Worker nodes? We have a few options:

1. Run the Hadoop balance utility. This is very slow and is primarily concerned with disk utilization—not data distribution. A nice feature is that it does guarantee that the data it's balancing will be continuously available for requests.

Invoking the Hadoop balancing utility is simple from the OS command-line.

[as the Hadoop super-user]:

```
$ hadoop balancer [-threshold <threshold>]
```

The optional threshold parameter enables you to specify how perfectly the disk utilization should even up across nodes. The default is within 10% - that is, the balancer will only work to relocate blocks until all nodes have a utilization factor within 10% of each other. The balancer can be interrupted at any time without bad consequence.

As a general rule, the best practice is to run the Hadoop balancer after adding new nodes to the cluster.

or

2. Remove and reload the SASHDAT data into Hadoop serially using SAS utilities. In testing, this is often much faster than the Hadoop balance utility and yields better distribution. But the data isn't 100% available in HDFS while this process is underway (due to save and rename). Afterward, the evenly distributed data must be reloaded to LASR.

This reload to HDFS can occur while the old, unbalanced table is in LASR. This ensures that users can still get to their reports and data (but not as fast or efficiently as before). Then during a planned outage window, drop the LASR table and reload it from the new SASHDAT table.

or

3. Proceed to load the unbalanced SASHDAT data into LASR. Then instruct LASR to rebalance the data for you across its Worker Nodes and save the data back down to Hadoop. The following SAS code shows an example:

```
/* Promote, Balance, and Save */

/* Setup a SASIOLA libref that specifies the LASR server */
/* where the unbalanced table resides */
libname example sasiola host="node0.example.com" port=10010
tag='sgf';

/* Use the data management capabilities of the IMSTAT */
/* procedure */
proc imstat immediate;

/* Specify the unbalanced table */
table example.the_table;

/* Print out the distribution stats so we can see */
/* how bad it really is */
distributioninfo;

/* Perform the balance - each node will be +/- 1 row */
/* A new temporary table (balanced!) is created */
balance;

/* Drop the original unbalanced table */
droptable;
```

```

/* Now reference the newly balanced temporary table */
table example.&_templast_;

/* Promote the temporary table to active status with the */
/* original table's name */
promote the_table;

/* Now reference the new table by its permanent name */
table example.the_table;

/* Print out the distribution stats for confirmation */
/* of balance */
distributioninfo;

/* Save the LASR table back down to SASHDAT on HDFS */
/* and replace the old, unbalanced table there */
save path="/path/in/hdfs" replace;

quit;

```

Promote, balance, and save is the fastest and best-balanced technique of the three discussed here. However, the data isn't 100% available while this process is underway (due to save and rename). But it's very close!

The time taken for a LASR drop to an in-memory table is often less than 1/1000th of a second. The time to promote the temporary table to active status can take as little as 1/100th of a second. That means that the total "outage" time for the in-memory table is far less than 1 second.

Chances are that your users won't even notice. Of course, it's prudent to schedule the balance-promote-save activity to a time when users are offline, but the necessary time is obviously minimal.

SITUATION REPORT

At this point of our recovery scenario, things are now looking a lot better:

1. While it was short a node, the Hadoop cluster created additional replica blocks of the data that was lost when the Data Node was down. Those new blocks were spread around the remaining nodes, taking availability and utilization into account. While still highly available, the HDFS blocks for SASHDAT were not as evenly distributed as they had been originally.
2. The physical rack server machine has its faulty power supply replaced. Its internal disk drives and their contents were not damaged, so the software services in support of Hadoop and LASR just needed starting up.
3. Hadoop was notified that the Data Node had rejoined the cluster. The Data Node was then able to immediately participate in any tasks as needed.
4. We restarted the SAS LASR Analytic Server so that it could pick up the extra Worker. We then reloaded the SASHDAT table into memory. The resurrected Worker likely has very little data of its own at this point and will have little to offer in terms of analytic calculations. The rest of the LASR Workers have an uneven distribution of data, which can further slow the perceived performance of LASR processing.
5. To correct the uneven distribution of data in HDFS and optimize LASR performance, we chose to lift the unevenly distributed SASHDAT table to LASR where we could then balance the records

across all nodes, promote the table to active service, and then save it back to HDFS for later use as needed.

CONCLUSION

After the unexpected hardware failure of the Data Node host machine, we were able to ensure continued availability of HDFS and LASR. When the host machine came back into service, we took the necessary steps to restore an even data distribution so that the SAS LASR Analytic Server could provide the best service and performance.

REFERENCES

The Apache Software Foundation. "HDFS Architecture Guide." Accessed February 2015. Available at https://hadoop.apache.org/docs/r1.2.1/hdfs_design.html.

The Apache Software Foundation. "Hadoop FAQ." Accessed February 2015. Available at <http://wiki.apache.org/hadoop/FAQ>.

Chansler, R., et al., 2013. *The Architecture of Open Source Applications* – "The Hadoop Distributed File System." Accessed February 2015. Available at <http://www.aosabook.org/en/hdfs.html>.

The Open Science Grid. "Hadoop Recovery." Accessed February 2015. Available at <https://twiki.grid.iu.edu/bin/view/Storage/HadoopRecovery>.

ACKNOWLEDGMENTS

The author would like to acknowledge the SAS LASR Analytic Server R&D team for their extensive contributions explaining the inner workings and design of the SAS LASR Analytic Server and the SASHDAT file format.

RECOMMENDED READING

- SAS® LASR™ Analytic Server 2.5: Reference Guide
- SAS® Visual Analytics 7.1: Administration Guide
- Apache Hadoop Commands Reference

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

Rob Collum

Sr. Technical Architect, Global Architecture & Technology Enablement
Professional Services Division, SAS Institute Inc.
(919) 531-0295
rob.collum@sas.com
www.sas.com

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are trademarks of their respective companies.