Paper 3046-2015

Thoughts on SAS® Visual analytics architecture for multiple customer groups

Jan Bigalke, Allianz Managed Operations & Services SE

ABSTRACT

SAS® Visual Analytics is a product to easily allow the interactive analysis of data. The product offers capabilities to analyze data with a visual approach. This paper discusses architecture options to configure a visual analytics installation that serves multiple customers in parallel. The overall objective is to create an environment that scales with the volume of data and also with the number of customer groups. This paper will explain several concepts to serve multiple customers and explain the pro and cons of each approach.

INTRODUCTION

SAS® Visual Analytics utilizes a web based approach to analyze data with the SAS® LASR™ Analytic Server. The web components follow the same approach as the SAS Intelligence Platform. This includes a SAS webserver and the SAS web application server. Together with the SAS Metadata Server and the SAS application servers, these components create what is called the "head node". The SAS® LASR™ Server hosts data in memory for the analytical workload. This In-Memory approach allows quick response times to the customer even with huge data volumes.

The SAS® LASR™ Analytic Server itself can run in two different modes:

The SMP (symmetric multiprocessing) mode, this is a single server with multiple processors and plentiful memory (RAM). The other option is the MPP (massive parallel processing) configuration. This is a distributed deployment model utilizing 4 or more servers, each with multiple processors and large memory allocation (≥ 256GB).

An environment for multiple customers has the requirement to scale with the customer demands. To best meet this requirement, the paper will focus solely on the MPP approach. Whilst the SMP approach allows a faster setup time, the MPP offers greater scalability through adding additional nodes to handle more data and more concurrent users.

For the SAS platform[1] there are some existing concepts published to provide a platform for multiple customers[2]. Also the approach to host multiple customers on a VA shared environment is already published [3,4]. These two papers introduce the concept of multi tenancy for the SAS platform and also for SAS® Visual Analytics. Based on these concepts this paper discusses the options to host several independent customer groups on the SAS® LASR™ server and the other SAS Visual Analytics components.

SAS PRINCIPAL ARCHITECTUE IN AN MPP ENVIRONMENT

The SAS Visual Analytics Installation and configuration guide[5] introduces the "SAS analytical cluster" as the distributed installation of the SAS LASR Server. The installation and configuration guide demonstrates some of the options to build such a cluster.
The picture on page 2 shows the LASR server Co-located with the Data Store. The data store is the Hadoop File system (HDFS), which is installed parallel to the LASR Server on every data node. The LASR server has its own Hadoop file format, called SAS HDAT, this format is optimized for fast loading of data into the LASR server. HDFS offers some key benefits[6]:

- Parallel I/O. The LASR Server can read data in parallel in high rates from a co-located data provider.
- Data redundancy: By default, two copies of data are stored in HDFS. If a server in the cluster fails, the LASR server instance on another server can read the data from a redundant block.
- Homogenous block distribution. That results in balanced memory utilization across the LASR server and reduces execution time.

[1] http://support.sas.com/documentation/cdl/en/biov/64876/PDF/default/biov.pdf
[2] http://support.sas.com/resources/papers/proceedings13/494-2013.pdf
[3] http://support.sas.com/resources/papers/proceedings14/SAS298-2014.pdf
[4] http://support.sas.com/resources/papers/proceedings14/SAS142-2014.pdf
[5] http://support.sas.com/documentation/onlinedoc/vs/index.html (SAS Visual Analytics 6.4: Installation and Configuration Guide)
[6] http://support.sas.com/documentation/cdl/en/inmsref/67306/PDF/default/inmsref.pdf

A program or analytic task is much more efficient if it's executed near the data it operates on. In this paper we will refer it as the "program near data" concept.

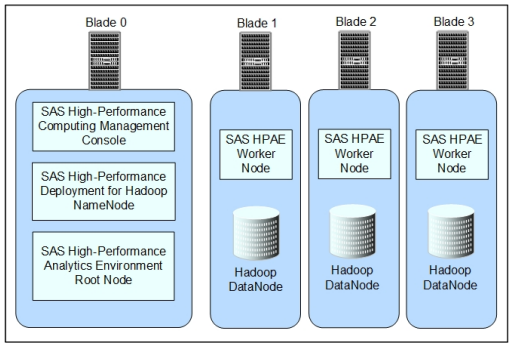

Picture 1 : SAS® analytical cluster [7]

Additional approaches also exist where the data is loaded directly from an external Hadoop System. This approach is called Parallel I/O: Co-Located Hadoop[8]. In this case only the LASR software is running on the data nodes.

Data nodes are typically built on physical hardware with lots of memory (≥ 256 GB RAM) in each server with local disks providing the data storage. From this available storage a co-located installation approach is implemented. In the concept of Parallel I/O: Co-Located Hadoop this available disk capacity is not used even if it's there. When using an existing Hadoop cluster, resource management for the Hadoop cluster is also required otherwise data loads will be slow, if the Hadoop cluster is already busy.

As an alternative to the default approach in the SAS documentation, this paper promotes an approach where the platform components are separated from the LASR Server. The configuration of the SAS environment follows the approach of a typical, multiple machine SAS deployment. The SAS 9.4 platform stack already supports scalability on all levels (Web Application Server Clustering and also new with 9.4, the clustering of the Metadata server). With this approach scalability and some failover support on the frontend (web interface) and also the SAS components e.g. for data loading is available.

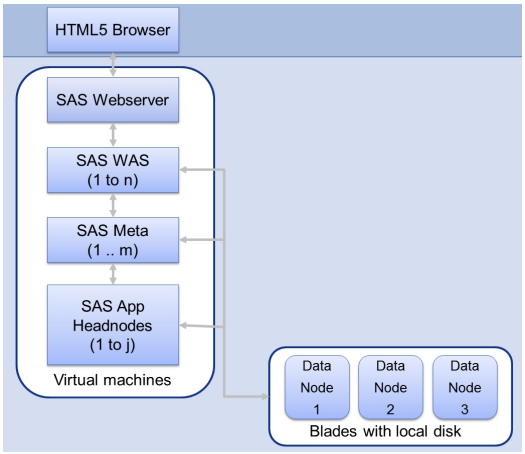

Picture 2 : head node on virtual machines

The implementation shown in the picture above utilizes virtual machines for the SAS components normally located on the head or root node. This concept can help in case of hardware failures for this server as the virtual machines can easily be migrated to other hardware. Even if you don't consider high availability with clustering, deploying the SAS installation on virtualized servers can help to reduce outage time if a hardware failure occurs.

[7] SAS® Visual Analytics 6.4 Installation and Configuration Guide [5], page 6
[8] http://support.sas.com/documentation/onlinedoc/vs/index.html (SAS Visual Analytics 6.4: Administration Guide)

WHY MULTIPLE CUSTOMER GROUPS IN ONE SAS VISUAL ANALYTIC ENVIRONMENT

SAS Visual Analytics has the potential to support several business lines with the provided analytical capabilities. The provided interfaces allow business users to do work that historically could only be done in IT departments. If the customer groups have the need to work independently and also requirements to protect the data and analytical results from each other the use case to serve multiple customers is already here. Advantages of a shared infrastructure compared to dedicated environment for each customer group are:

- Time to market. Setup of a complete dedicated environment for a customer is time consuming, from the hardware order to the installation and Software configuration.
- Efficiency: If you have several customer groups but each has only a small number of users the resource overhead in web application server is huge (baseline of resources to bring the web applications online).
- Promotes the option to start small (small start invest) and grow later more easily form the customer perspective.

The Challenges for a multi customer approach are:

How to configure the SAS Application Server and also the HDFS & LASR-Server whilst considering scalability, performance & cost? Also the capabilities of the GUI components of SAS Visual Analytics to handle this approach needs to be validated.

MULTIPLE CUSTOMER GROUPS IN ONE SAS VISUAL ANALYTICS ENVIRONMENT

The SAS Visual Analytics product as supplied to the customer today has no built in multi tenancy functionality. To achieve full multi tenancy on all levels is, from the status of today, it is not possible based on the provided capabilities of the product.

However if the customers group's consist of different departments, and are willing to live with a not complete multi tenancy implementation, security options on the operating system and within the SAS system can support several customers inside one Visual Analytics environment. Some configurations in this approach can't be adjusted per customer like the web tier configuration and have to be agreed form every user group.

The approaches to host multiple customer groups are typically based on multiple SAS application server contexts in the SAS compute tier. Application security inside the SAS metadata server protects the application server contexts from each other on the metadata level. Each customer group can only see and use their assigned SAS application server context. On the file system side the separation of communities is achieved through operating system security - primarily with UNIX groups and independent technical accounts like the SAS sassrv user.

The following chapter discusses options to extend the SAS platform multi tenancy concepts to the SAS LASR components in more detail.

Configuration option 1 is shown in the figure below. Located on the head node are the SAS server contexts for each customer. Every customer has an independent area where they can store their input data for the LASR servers (called Input data storage in the picture). Parallel to the SAS server contexts are the LASR server configurations for each of these SAS server contexts. On a data node there are dedicated LASR processes for every customer group. The LASR process for each customer has an assigned amount of memory limited via CGroups on the operating system layer on every node. Configuration option 1 utilizes one shared Hadoop file system (HDFS) instance for all customers. On the head node one Hadoop Name Node (HDFS NN) is configured. HDFS stores the data in blocks distributed over each of the available nodes with redundancy. The picture visualizes this in the File Distribution layer. For example on every node blocks from "File 1" are available. The HDFS NN determines the mapping of blocks to the HDFS data nodes.

VA 6.4 configuration option 1

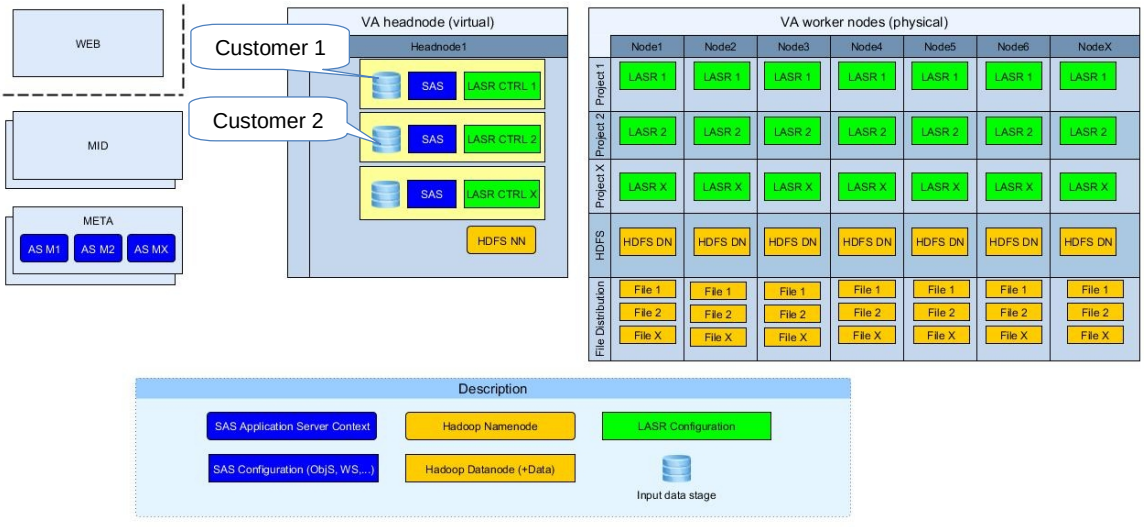

For a limited number of customers this approach is feasible, however with larger customer volumes the resource overhead to run the high number of LASR instances is extremely high. In addition, this approach does not recognize different capacity requirements between the customers. Multi tenancy is also necessary on the HDFS level to fully support security and separation aspects on data storage layer. Multi tenancy in HDFS (Hadoop) is possible but enforces the usage of the Kerberos authentication protocol[9]. This authentication protocol is not supported for VA 6.4 from the LASR to HDFS. Lastly the scalability on the head node is ultimately limited to the size of the server, because today the HDFS NN must be collocated with the LASR Head Node configuration.

In configuration Option 2 the LASR is not started on all nodes for every customer. This configuration allows support for the different in-memory data capacity needs of each customer. Based on the required capacity for memory, the required numbers of LASR processes are started on a subset of the total infrastructure. This enhancement helps to reduce the resource consumption for the application infrastructure on the data nodes as overall, fewer LASR instances are running across the cluster.

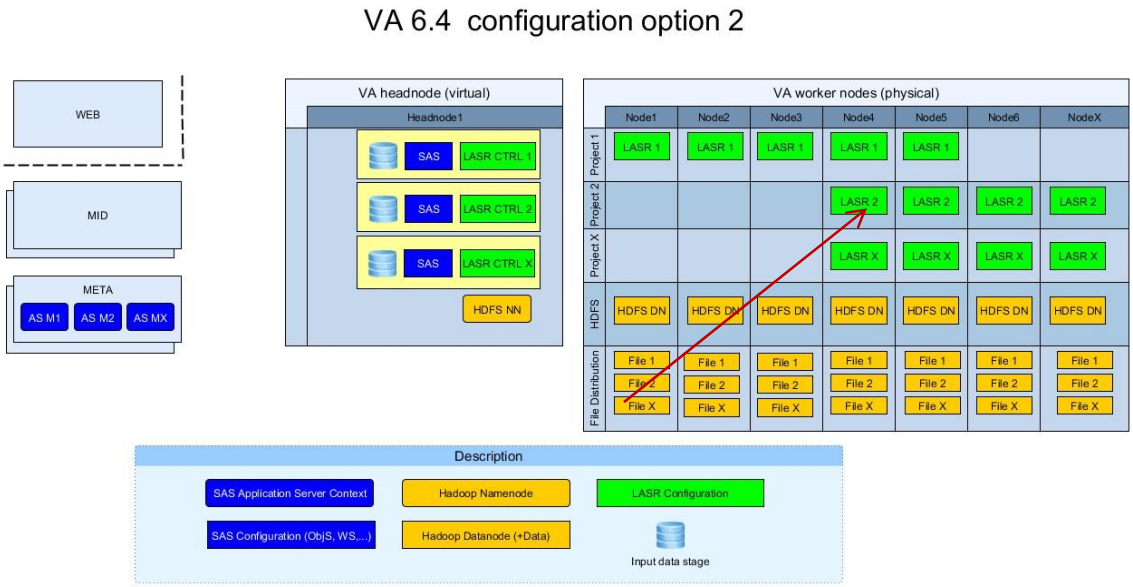

Rating: Resource efficiency is improved by running fewer LASR server instances, however the memory/data mapping is broken and data has to be transferred over node boundaries. Over and above the limitations of Option 1, this option is handicapped by impact to LASR's use of the local storage with the proprietary SASHDAT format. The parallel load of the data can't be done on the local server anymore and use also network communication.

Configuration Option 3

This approach utilizes a separate HDFS instance with a separate HDFS Name Node per customer. Each HDFS instance use dedicated ports, configuration files and quotas to manage the file system space.

The data nodes for each customer can be distributed across the data node cluster. Each configuration uses only the capacity needed to support the requirements of each customer. This configuration provides horizontal scalability of the head node, whereby a head node instance can be added to the existing configuration if resources become insufficient on the existing head node.

Rating: On the data nodes the data is co-located to the LASR processes. The HDFS file system is only assigned to nodes where the customer has LASR instances started. All the introduced benefits: Parallel I/O, data redundancy and homogenous block distribution for efficient execution are in place. The head node can scale with the customers. This approach scales with data (additional data nodes) and customer groups (new head node instances and new LASR instances based on the needed capacity). The concept provides every customer inside one physical cluster their own individual cluster configuration, utilizing only a subset of the overall capacity of the cluster.

[9] http://hortonworks.com/wp-content/uploads/2011/10/security-design_withCover-1.pdf

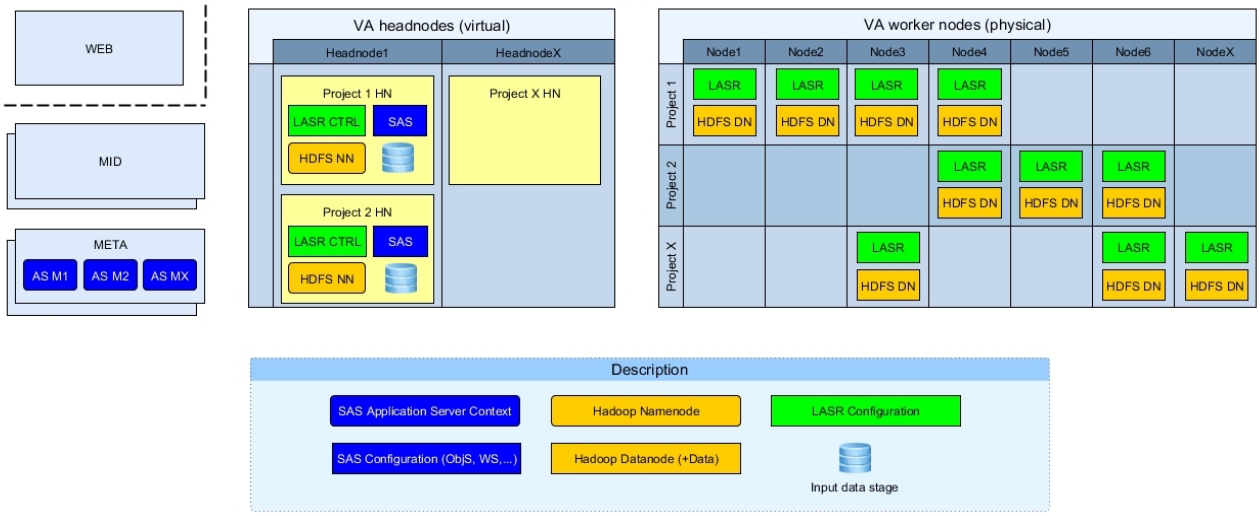

The disadvantage is that the duplication of the HDFS Name Node and some of the components like the HDFS monitor cannot handle multiple HDFS NN. Therefore these components cannot be surfaced to every customer and reduce the available capabilities of the SAS Visual Analytics product for individual customer. The analytical capabilities are all available to the customers so that the limitations are small compared to the achieved benefits.

RESSOURCE MANAGEMENT

SAS Visual Analytics is not designed for multiple instances of the LASR Server on a single server. If several customer groups use one SAS Environment together resource management is important. Resource control should ensure that every customer group gets their ordered share of the overall capacity of the environment. The resource management is focused on the head node and data node. The SAS Web Application Server and the SAS Metadata are totally shared so a resource management cannot be controlled on the infrastructure layer.
The discussed Visual Analytics infrastructure is built on the Linux operating system. In Linux the CGroups concept provides capabilities to manage hardware resource usage. The paper form the last SGF 2014 shares some thoughts[10] on the use of CGroups.

For the Visual Analytics use case the focus is on CPU and memory resources. With CGroups and defined CPU resources for all customers the compute capacity for each customer is controllable. For the management of the memory usage also CGroups can be used. However, the usage of memory resource control has some limitations. The LASR process does not recognize the memory limitation via CGroups and will fail if the LASR Server tries to allocate more than the available memory. Also the display of the free capacity within Visual Analytics administration does not work anymore, as the free capacity relates to capacity on a data node not the assigned capacity via CGroups.

The LASR® Analytic Server 2.4[11], included in SAS® Visual Analytics 7.1, additional capabilities were introduced to manage resources. For CPU usage also CGroups are used. Two settings are related to memory usage. TKMPI_ULIMIT set a memory limit for the entire LASR process and TKMPI_MEMSIZE a limit to the size of in-memory tables and memory that is used for processing actions.
For the HDFS file system quotas can be used to limit the usage[12].

FURTHER CONSIDERATIONS FOR A SIMPLYFIED DEPLOYMENT

In the currently available versions of SAS Visual Analytics, some tasks will use the SAS workspace server to execute SAS code. In an MPP configuration the commands to the remote components on the data nodes are executed with remote SSH calls. Remote SSH calls require the creation and deployment of SSH keys.
If the Workspace Server is executed with the credentials of the end user of the web interface, each user that has the capabilities to execute these tasks needs their own SSH keys and access to each of the servers. To avoid the need of granting access to every data node and the creation of SSH keys the workspace server can be configured with token authentication. In this case all tasks are executed with a technical account defined for these tasks.

[10] http://support.sas.com/resources/papers/proceedings14/SAS289-2014.pdf
[11] http://support.sas.com/documentation/cdl/en/inmsref/67597/PDF/default/inmsref.pdf (page 435)
[12] http://hadoop.apache.org/docs/stable/hadoop-project-dist/hadoop-hdfs/HdfsQuotaAdminGuide.html

CONCLUSION

Utilizing SAS® Visual Analytics in an environment with multiple customer groups is possible configuration. The most flexible approach utilizing a Hadoop Name Node for each project has some impacts on the functionality of the SAS software, because from the design only one Name Node is supported. If these limitations are not critical, an infrastructure that scales with the number of customer groups and also the volume of data can be built. An improved support from SAS for configurations like those discussed earlier would be helpful - for example components that can handle LASR configurations based on the used SAS Application server context.

The reader should be aware that the effort to implement this concept is significant. First it is necessary to setup the platform components in a multi customer configuration and in the next step to adjust the LASR configuration. For the components new ports and configurations have to be created. The demand for Visual Analytics should be there for several different customer groups otherwise the effort and efficiency gains will not be realized.

REFERENCES

SAS® 9.4 Intelligence Platform
http://support.sas.com/documentation/cdl/en/biov/64876/PDF/default/biov.pdf

TT in the sky: Building a private SAS® cloud
http://support.sas.com/resources/papers/proceedings13/494-2013.pdf

SAS® Visual Analytics for the Three Cs: Cloud, Consumerization, and Collaboration
http://support.sas.com/resources/papers/proceedings14/SAS298-2014.pdf

Security Scenario for SAS® Visual Analytics
http://support.sas.com/resources/papers/proceedings14/SAS142-2014.pdf

SAS® Visual Analytics 6.4: Installation and Configuration Guide
http://support.sas.com/documentation/onlinedoc/vs/index.html

SAS® Visual Analytics 6.4 Administration Guide
http://support.sas.com/documentation/onlinedoc/vs/index.html

Adding Security to Apache Hadoop
http://hortonworks.com/wp-content/uploads/2011/10/security-design_withCover-1.pdf

SAS® LASR™ Analytic Server 2.3 Reference Guide
http://support.sas.com/documentation/cdl/en/inmsref/67306/PDF/default/inmsref.pdf

SAS® Grid Manager, SAS® Visual Analytics, and SAS® High-Performance Analytics: Sharing Hardware and More
http://support.sas.com/resources/papers/proceedings14/SAS289-2014.pdf

SAS® LASR™ Analytic Server 2.4 Reference Guide
http://support.sas.com/documentation/cdl/en/inmsref/67597/PDF/default/inmsref.pdf

HDFS Quota Guide
http://hadoop.apache.org/docs/stable/hadoop-project-dist/hadoop-hdfs/HdfsQuotaAdminGuide.html

ACKNOWLEDGMENTS

The author gratefully acknowledges Andrew MacFarlane of SAS Institute for reviewing the content of this paper. Thanks.

CONTACT INFORMATION

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

Name: Jan Bigalke
Allianz Managed Operations & Services SE
jan.bigalke@allianz.com