

## How many Genies can you fit in one bottle? – Multi Tenancy and all the wishes you want

Ross Sansom, RBS;

### ABSTRACT

With data becoming bigger and bigger year after year, it is important to have the right tool set and platforms to meet the need of the organization. The idea of one platform to meet all of your needs is something we all think would be the perfect solution. Somewhere to bring all your business needs and thereby make the management of the platform a lot easier. Royal Bank of Scotland (RBS) is now three years down the line, and the dream has become the reality. This paper looks at how RBS implemented a multi-tenancy platform, which is internally known as Genie, and what was discovered along the three-year journey. This paper covers: building tenancies and how using one format to fit all can work; user management; data segregation, or how many SAS® LASR™ Analytic Servers do you really need?; security and using the same model to meet all requirements. One SAS® 9.4 platform, 10,000 users, and 14 tenancies; did the Genie help build our dream...

### INTRODUCTION

Data is becoming more and more important each day to organisations and RBS is no different. With data being so important, it was also fundamental to having the correct toolset and set-up. RBS know this and in 2015 ran a proof-of-concept of SAS® Visual Analytics and its uses to the bank. The POC was a success and this led to the building of Genie, the bank's SAS® 9.4 Strategic Shared Grid. The intention was for Genie to modernise a majority of the existing SAS® estate by replacing siloed use of SAS® with a central platform.

The bank had a dream for Genie to be a multi-department platform that met the needs of analyst community across the two biggest analyst teams in the bank. The bank worked with SAS® to start building its two-tenancy platform with access to fifteen different SAS® products. It didn't take long before more and more business areas wanted to be a part of the Genie dream.

This paper will look at the building of the Genie dream and how we went about on-boarding multiple business areas with the same design and the power of the Little Green Button bringing automation in to the build process. It will also look at the user management aspect as well as data segregation.

### WHAT IS GENIE MADE UP OF?

At the time that Genie was designed, it was in the running to be one of the biggest SAS® Visual Analytics platforms in the world; consisted of the following:

- A 9.4 SAS(R) Grid which consists of 59 physical LASR servers, 7 Grid servers and 3 Metadata/ Mid-tier servers. Giving 30TB of in-memory data storage with another 40TB of HDFS Storage which is collocated on the same servers as LASR.
- The main SAS® products being SAS® Visual Analytics and SAS® Enterprise guide but there is also an additional 8 SAS® Products as a part of the platform.
- The ability to connect to a number of different data sources including Hadoop, Teradata and Main Frame.

- Genie was sized for two business areas, 9,000 report consumers and 1,000 analysts.

## WHAT DOES MULTI-TENANCY REALLY MEAN?

“Multitenancy is a reference to the mode of operation of software where multiple independent instances of one or multiple applications operate in a shared environment. The instances (tenants) are logically isolated, but physically integrated.” - Gartner IT Glossary

At RBS, in the context of Genie, multi-tenancy means to have a shared platform for multiple business areas allowing segregated and collaborative work in one place. Although RBS went with this approach but there were a number of different approaches considered. These approaches are discussed in another SAS® paper “Multi-Tenancy in SAS® - Is It Worth the Fight?” (Blake et al., 2016).

## MAKING ONE DESIGN THAT FITS ALL

With Genie becoming more popular around the bank it was important that there was a way of building new tenancies quickly and effectively as well as being able to add new secure areas within the tenancy which we call channels. This is why RBS came up with a design of tenancies and the distribution channels within them that would allow for segregated work within the tenancy. The types of channels that RBS has will be covered later on in the paper. To stop the building process becoming overly complicated it was decided that RBS would have a design that would fit the requirement of all users of Genie which lead to the vanilla build.

### THE VANILLA BUILD

When it comes to on-boarding it was decided that there would be a standard set up for all new business areas coming on to Genie. This was named The Vanilla Build which could be customized as time goes on to meet the business requirements.

The vanilla tenancy build consists of:

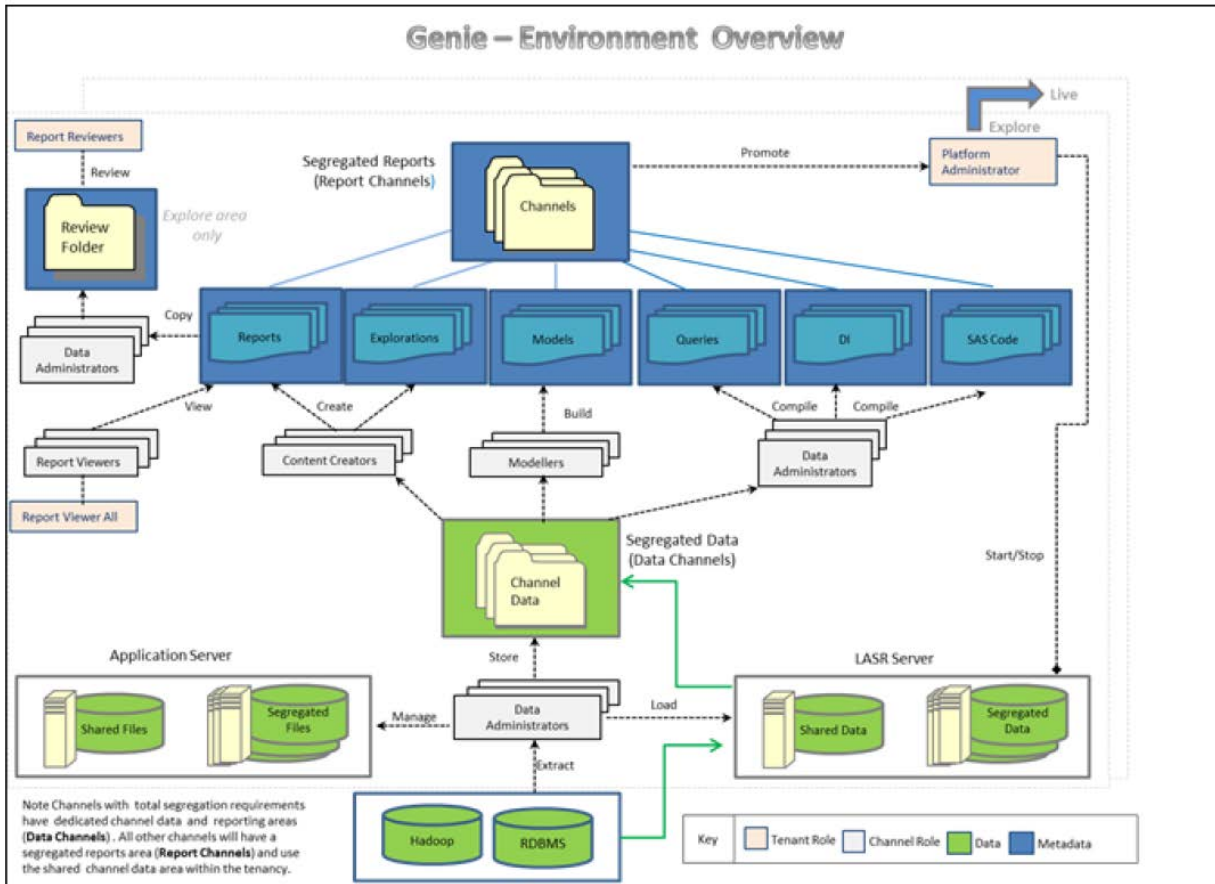
- A shared folder structure.
- 4 role groups with different permissions.
- Access to the different tools depending on role group.
- 1 shared application server per tenancy.
- 1 shared LASR server definition per tenancy.
- 1 reporting channel (unless the first channel in the tenancy has been requested as a dedicated data channel).

### ***Types of channels:***

The business area would then have a choice of two different types depending on their security needs.

*Reporting channel* – this type of channel is secured with its own access groups and uses the shared tenancy application server, HDFS and LASR server for processing and storing of data. This type of channel should be used for reports that uses data that is used by the tenancy and is stored on the tenancy shared LASR server. This type of channel would be used for less sensitive projects.

*Dedicated data channel* – this type of channel was designed for a more secure area on Genie. This has the basics of a reporting channel with added security around the data side. The dedicated data channel gets its own LASR, HDFS and Application server but it also has access to the tenancy shared LASR and Application server so analysts can bring in data that is being used by the wider tenancy as well. Figure 1 below show how this works. Data is brought in from the data source and loaded in to the segregated LASR server which means it can only be access' by people within the channels security groups and not accessed by everyone in the tenancy similar to a report channel. The SAS® Visual Analytics reports are then built from that data and stored in the secure folder structure which again can only be accessed by users in the security groups for that channel.



**Figure 1. Metadata Design for Genie (Blake et al., 2016)**

Each channel type comes with four pre-defined groups which control access to folders as well as tools.

Once a tenancy is up and running the business has the choice to customize their tenancy with more channels and get the choice of two different types of channels. This means they can expand their tenancy over time to meet the needs of their area. The Genie support team continuously checks with the business to see if they need more channels built or if there is any channel are no longer being needed and can be deleted. To this date every business area has gone with the design which has been spoken about above and if the business was to want a more custom builds this would be very complex and take some time

to think it though. This would also mean we would either have to do it by hand or make a big change to the Little Green Button to do it.

## THE GREEN BUTTON THAT CHANGED EVERYTHING

With the Genie platform becoming popular around the organisation it became apparent that there was a need to think of a way to make building tenancies and channels easier. This led to a piece of joint work with SAS® called the Little Green Button to help bring the multi-tenancy dream to life even faster.

The Little Green Button actually came out of a bigger project called the Big Red Button which was designed to completely automate the build process as shown in Figure 2 below. The idea was there would be no human intervention and almost the end to end process would be completely automated. The only parts not automated would be things that needed to be done on external systems. This was the dream but due to time and how complex it was the red button was never complete but this did make way for the green button which has allowed us to automate most of the metadata tasks.

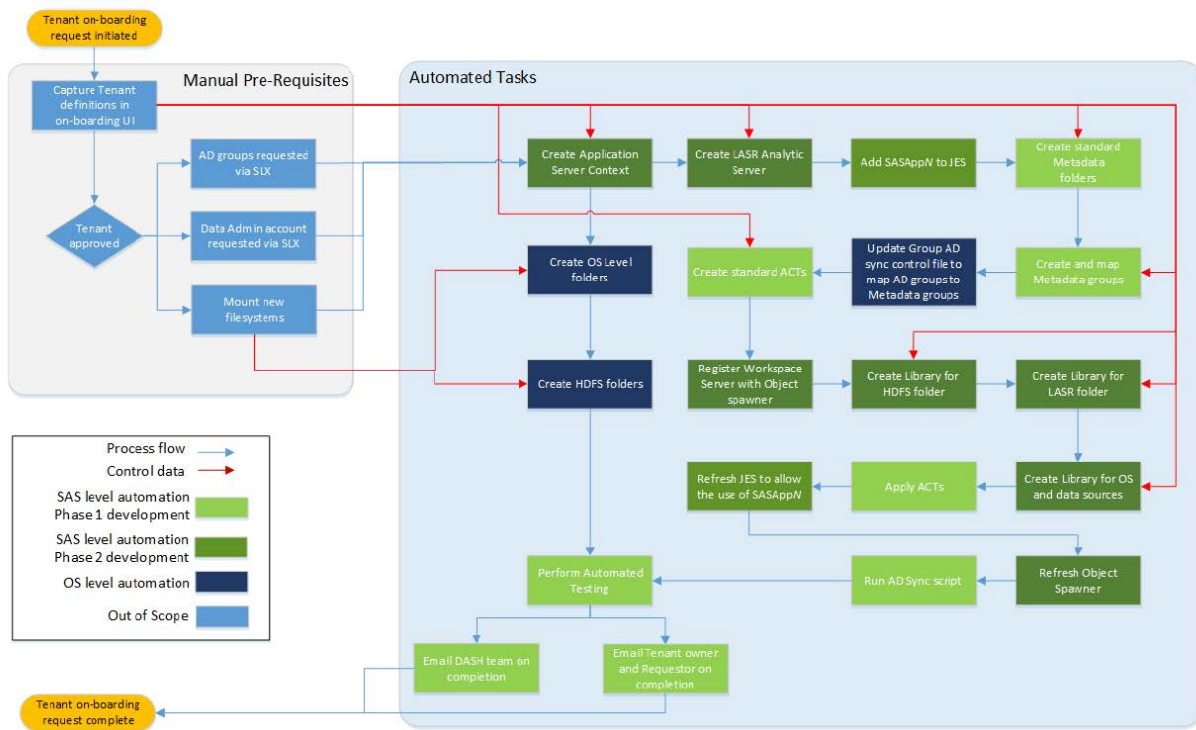
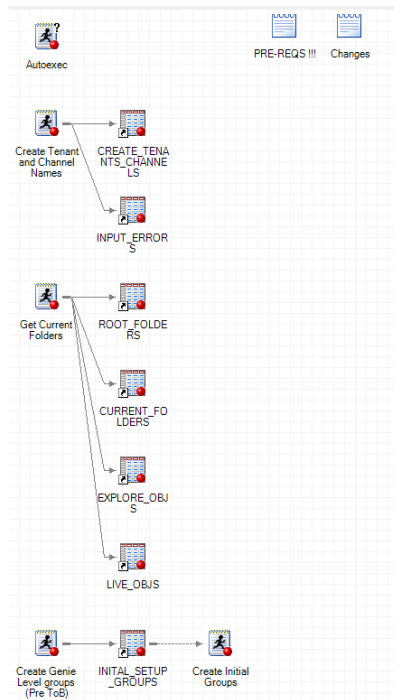


Figure 2. Red Button End-to-End Tenant On-boarding Process (Blake et al., 2016)

## What is the Little Green Button?

The Little Green Button is a SAS® Enterprise Guide project designed to make the building process of metadata as easy as possible allowing more time to concentrate on the complex part of the build. The Little Green Button brought the time it took to build a single reporting channel down from two days' work to just five minutes. It also reduced a new tenancy or dedicated data channel build from two weeks work to a couple of days. The figure below shows just the first stage of the Little Green Button.



**Figure 3. Part of the Little Green Button**

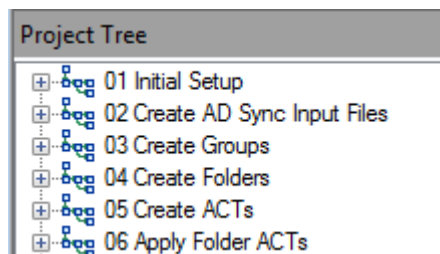
The Little Green Button does the following:

- Builds the folder structures
- Builds the access control templates
- Secures the folders with the access control templates
- Builds all of the role groups
- Adds the role active directory sync

All of tasks manually would of take hours or even days' worth of work. There is currently no plan to extend what the Little Green Button does but this might be something RBS looks at in the future.

### How does the Little Green Button work?

Once the Little Green Button is open an autoexec runs to check the environment you are connected to. This ensures you are running the code on the correct environment if you have more than one environment like an NFT environment.



Now the autoexec is done all that is needed to be done is press the run project button and let it do its thing. The Little Green Button then works through the following stages:

1. Read in the input for the new channels or tenancy. This is in a form of a data step. The example below shows the format of the data step that we would put in to the Little Green Button. This is the only manual part of the Little Green Button but is also the most important part to get right. This needs to be checked a number of times to make sure there is no spelling mistakes and that it is in the correct format that the Little Green Button will recognize.

Example:

```
Tenancy Name,Channel Name,Tenancy short code,channel type indicator  
Test,Genie Channel,TESTGEN,R (R= Reporting Channel D= Dedicated Data  
channel)
```

This input would create a report channel in the tenancy Test called Genie Channel.

2. Check to see if any of the elements already exist including tenancy names, channels and access control. This is a safe guard that ensures that no tenancies or channel can be named the same. This is so to make sure that nothing is created twice when it shouldn't be.
3. Start the build process by creating the following:
  - a. High level groups – these groups are the groups that provide the capabilities to the users. These would have been created on the first run of the Little Green Button but is re-run just in case there has been any new groups defined or the existing groups have been damaged in any way.
  - b. Creation of the tenancy / channel template.
  - c. Creation of the Access control template (ACT).
  - d. Creation of the Folder structure and application of the ACT Template.
4. Little Green Button then moves on to updating Active Directory Sync. It does this by updating two files stored on the UNIX server as a part of Active Directory sync. These files are group\_info.csv which stores the information about each of the groups and group.txt which holds a list of all the groups on Genie. This can be done manually and would take about 5 minutes to update these documents where it only takes Little Green Button a couple of seconds and it means there is no risk of a mistake as it feeds off the date input which has already been checked.
5. The next thing to be built is the Tenancy/ Channel groups. This is where the groups for the different roles are created and the capabilities and permissions are set using the pre-defined template. If RBS was going to build this manually this would take about 5- 10 minutes per group where as Little Green Button does a least 4 in under a minute. By automating this it means that there is no chance of anything going wrong in the setup of the groups and no chance of human error.

6. After the groups have been created the next thing to be created by the Little Green Button is the folder structures. The folder structure has many layers to it depending on channel type and if the tenancy even exists. The folder structure is designed to cover all of the scenarios. It also means that the folder structure works with all of the tools and everything is in one place. This allows for better collaboration across the teams. This would be a really heavy and repetitive process which could take hours of work per tenancy or channel folder structure.
7. The second to last thing to be created is access control templates. This is where the groups are added to the access control and the permissions are set for that template. This is one of the most important tasks to make sure it gets set up correctly. This is why having an automated solutions is so great because it doesn't allow for any issue with the setup of a fundamental part of the security design. Doing this manually could lead to security issues on the platform for instance if the person setting the access control template up is not concentrating and does the set up wrong. Setting up a template can be a very timely task with 1 access control template taking 10 minutes to set up. The Little Green Button does this in minutes saving hours of work and ensuring the integrity of the security design.
8. The final thing that Little Green Button does is the most important. This is where the ACT's are applied to the folder structure to secure them.

The Little Green Button has really saved time but there are still a number of manual tasks to be done after the Little Green Button has run if the team is building a tenancy or dedicated data channel like building of libraries and configuring applications servers.

## USER MANAGEMENT

With a platform sized over 10,000 different users managing them all was never going to be an easy task. To make the on-boarding process as easy as possible for users it was decided to use the bank standard user authorization system and active directory to sync users in to groups. The Little Green Button supports the process by adding any new role groups to the active directory sync so that the banks authorization system and Genie syncs together.

The user will go through the following sets to get access to Genie:

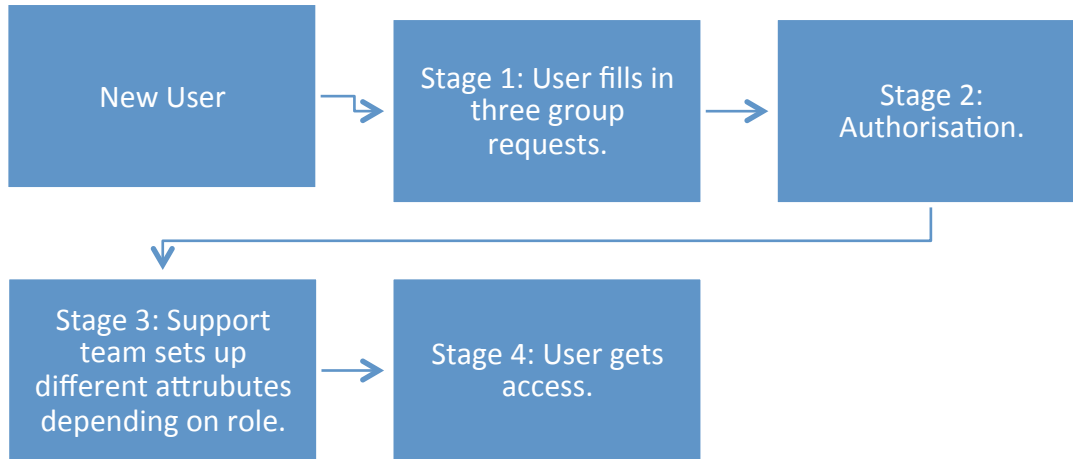
Stage one: The user will fill out three access requests to get access to Genie. These groups allow access to different parts of the system. The first group is for allowing basic access level to Genie and is a group that all users must be in. The second group is to allow access to the UNIX area attached to the tenancy or dedicated data channel. The third and final group is for the role/level of access they need with in the tenancy or channel they need access to.

Stage two: Authorization which is done by different people. For the first two groups then it will be authorized by the Genie support team. The role/level of access group will be authorized by the local authorizers in the business areas.

Stage three: The Genie support team use a in house built SAS® Enterprise Guide Project call the gap analyzer to check the users profile to see what attributes are needed. These are things like does the User need a UID, Group ID or Home folders added to their UNIX profile

and shows if they are needed for the level of access being granted. It also checks to make sure that all three groups have been authorized.

Stage four: User's access is granted.



**Figure 4 End to end Genie user on-boarding**

The longest part of this process is stage three as this can need other teams to make changes to a user profile. If the user already has all of the attributes present then it means they can be on-boarded faster. It normally takes between 24 hours and 48 hours to get a user on-boarded from end to end. User management has been one of the biggest issues of Genie and one recommendation to anyone about to set up a platform of their own is about your joiner, mover and leaver's process make sure you really think it through at the start as it may be a big head ache at the end.

## DATA SEGREGATION

With more business areas on Genie it's important that only the correct people have access to the data that they should. This is why data segregation was really important in the Genie design. The main segregation being between the different tenancies and this was done by providing each tenancy with its own LASR and HDFS area as well as their own application server mean there was no way of any data being used by anyone who it should not be used by. This is also backed up by the security groups and access control templates.

There is also segregation available in each of the tenancies via the dedicated data channel. This again a lead to tenancy having more LASR servers, HDFS areas and applications servers being set up. Every time a new tenancy or dedicated data channels is requested it means another LASR server, HDFS area and dedicated data server. This has led to Genie having 150+ LASR servers and the same in application servers which is a massive over head for the Genie support team.

RBS has put a number of things in place to ensure they know what is going on with all of the LASR servers to insure and the best service for all of the users on Genie. This is done in a number of ways. RBS uses the power of SAS® Environment Manager to set up alerting that notifies the support team by e-mail if the LASR servers stop all of a sudden, if the CPU gets too high or even if the space on the servers is full up. This allows the team to respond to issues as soon as possible. The other way is some in house scripts that run's every morning a 5am that does a full check on application servers and LASR servers to make sure they are up and running and then provides a report on any issues. Having this functionality



in place has meant that Genie has a high available percentage and allows the team to be proactive in the support.

## WHERE WE ARE NOW

The Genie platform is now the biggest it has ever been with 15 different tenancies and over 14,000 users and growing every day. The platform is now home to more than 3,000 business reports. Genie delivered on its dream of a multi-tenancy platform that meets the needs of the business. It is not providing real business benefits and is allowing us as an organization to really help our internal and external get the best service they can.

## CONCLUSION

In conclusion this report has shown that multi-tenancy can work if it is thought through correctly. The Little Green Button really did save time and make the whole build process a lot simpler especially with the one design fits all. If it was not for the vanilla build that really does work for everyone it would mean that the Little Green Button could not exist. Genie is now bigger and better than it has ever been and is showing real benefit to the bank and the bank's customers.

## REFERENCES

Blake, C., Iglov, S., Pournami, D. and Goienola, E. (2016). Multi-Tenancy in SAS® - Is It Worth the Fight?. [epaper] pp.1,10,11. Available at: <https://support.sas.com/resources/papers/proceedings16/11684-2016.pdf> [Accessed 17 Mar. 2019].

Gartner IT Glossary. (2016). *Multitenancy*. Available: <https://www.gartner.com/it-glossary/multitenancy>. Last accessed 15/03/2019.

## ACKNOWLEDGMENTS

The author would like to thank Ekaitz Goienola, Robin Thomson, Andrew Dent and Tony Parnell for their support and Jonathan Lloyd the man behind the final green button which really changed the build process for RBS.

## RECOMMENDED READING

- SAS® Paper - Multi-Tenancy in SAS® - Is It Worth the Fight? - <https://support.sas.com/resources/papers/proceedings16/11684-2016.pdf>
- SAS® 9.4 Intelligence Platform: System Administration Guide, Fourth Edition
- SAS® 9.4 Intelligence Platform: Security Administration Guide, Second Edition
- SAS® Environment Manager 2.5 Administration: User's Guide

## CONTACT INFORMATION

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

Ross Sansom  
RBS  
+44 (0) 1316261361  
Ross.sansom@rbs.co.uk  
LinkedIn: [www.linkedin.com/in/ross-sansom](http://www.linkedin.com/in/ross-sansom)