

All In: Integrated Enterprise-Wide Analytics and Reporting with SAS® Visual Analytics and SAS® Business intelligence

Jonathan Pletzke, University of North Carolina at Chapel Hill; David Franklin, SAS Institute Inc.

ABSTRACT

In 2013, the University of North Carolina at Chapel Hill (UNC) initiated enterprise-wide use of SAS® solutions for reporting and data transformations. Just over one year later, the initial rollout was scheduled to go live to an audience of 5,500 users as part of an adoption of PeopleSoft ERP for Finance, Human Resources, Payroll, and Student systems. SAS® Visual Analytics was used for primary report delivery as an embedded resource within the UNC Infoporte, an existing portal.

UNC made the date. With the SAS® solutions, UNC delivered the data warehouse and initial reports on the same day the ERP systems went live. After the success of the initial launch, UNC continues to develop and evolve the solution with additional technologies, data, and reports.

The intent of this presentation is to touch on a few various elements required for a medium to large size organization to integrate SAS® solutions such as SAS® Visual Analytics within their infrastructure.

INTRODUCTION

SAS Visual Analytics, along with other SAS BI tools, has become a standard, integrated platform within the organization at UNC.

In order to accomplish such a large rollout with an aggressive time schedule, much integration had to be performed.

Multiple environments were set up for both SAS BI and SAS Visual Analytics, including demo, development, test, and production. These environments were integrated with existing enterprise assets such as network devices, directories, authentication, firewalls, and so on. All this was done while the staff was coming up to speed on the latest server-based SAS® solutions, while development of the ETL was ongoing, and in the midst of a massive ERP implementation.

This paper will explore a few of those points of integration and cover the considerations and lessons learned during the initial implementation.

GOAL: NOT SAS FOR SAS' SAKE, BUT ENTERPRISE WIDE ANALYTICS AND REPORTING

SAS has been used at UNC for many years, both as single user desktop SAS, a distributed research computing platform, and in a couple of SAS/IntrNet configurations. This use covered areas of research and administration. The first UNC data warehouse for HR, Payroll, and Finance was built on SAS/IntrNet starting in 2002, and was frozen when the new ERP and data warehouse went live in October of 2014, and replaced by a new integrated enterprise wide solution that is running ETL and reports in SAS Visual Analytics.

As UNC started implementing PeopleSoft, the need arose to provide a reporting and analysis solution on the administrative data that comes out of the newly implemented PeopleSoft applications, including the PeopleSoft EPM data warehouse. In addition, other reporting and analysis would be required from other data sources and related enterprise systems. In prior years, much effort was spent trying to get the PeopleSoft data warehouse and reporting working with tools from multiple other vendors. After many years of failed attempts to satisfy the reporting gaps, UNC turned again to SAS to pursue an enterprise wide analytics and reporting solution and built a new data warehouse from scratch.

STRATEGY: VISUAL ANALYTICS FOR REPORTING, DI STUDIO FOR ETL, WEB REPORT STUDIO IF NECESSARY

In December 2013, representatives from the UNC technical and business teams met with SAS resources to discuss overall needs and a plan of attack for execution. With an arsenal of SAS Solutions at their disposal, it was recommended that UNC utilize SAS VA for reporting on the large amounts of data, and connect to the disparate data sources and develop the necessary data transformations using SAS Data Integration Studio. As a supplemental report delivery, SAS Web Report Studio could also be used to fill any other needs.

UNC learned that when integrating a SAS Solution, you must consider the hardware and architecture necessary to achieve success. In order to determine the needed hardware SAS offers a free sizing exercise through the SAS Enterprise Excellence Center (EEC) to take your expected SAS usage and translate it to an expected resource need. The sizing exercise is only as successful and useful as the information provided so it is essential to have the following:

1. Know your business process and your business need. For what will you use SAS? How are you going to use SAS, who is going to use SAS and for what reason?
2. Know your data sources and translate the sum of all data, to the actual summarized or usable data, including pre-joined reporting tables for in-memory applications. Not all data is created equal, but it might be important. Be sure to consider your data and its size as this can greatly impact your hardware needs. Are there major data manipulations that will be needed to bring your data to a consumable state? What is a reasonable timeframe to meet your needs (i.e. Processing 100 GB of data in 5 seconds requires much more hardware than processing 100 GB of data in 5 hours)?
3. Know your user base and their function. Create a “profile” of each user role. Understanding your users and their function will help determine permissions and groupings in your SAS solution.
4. Know what environments are needed for your success. What is your development life cycle? Do you need a separate Development, Testing, and Production environment? Are you really going to be OK with a single environment? Is a disaster recovery instance needed to meet your end user needs or your customer expectations? Please note, the answers to these questions may impact your SAS licensing structure.
5. Be flexible with your expectations. Start with the end in mind, and the path to that end can change at any moment. New business needs arise and SAS can scale to fit your needs. Keep an open mind to utilize resources. You may find you need to move some chess pieces around to better fit your requirements.

UNC initially decided to separate the ETL processing from the report development and put them in separate environments. Further separation would be realized as it was also decided to have the environments on heterogeneous operating systems. For ETL work, Windows was chosen. For reporting, RedHat Linux was chosen. This approach was to help divide and conquer tasks and utilize available resources for each platform. In addition, the business requirements for reporting – being large amounts of data for a broad audience – a distributed instance of SAS Visual Analytics was essential and with that came the RedHat Linux requirement.

After going through the EEC sizing exercise, UNC began securing and configuring the hardware topology for their initial enterprise wide analytics and reporting solution rollout.

DATA: FINANCE, HR, PAYROLL, STUDENT, ADVANCEMENT, AND MORE

The data to be incorporated is broad and deep. PeopleSoft includes tens of thousands of tables, of which we are primarily interested in only a few thousand. The source data spans three different PeopleSoft database infrastructures, as well as the companion systems required for operation of a large research and education enterprise. A diagram showing the input databases, the SAS Data Integration ETL processing that delivers data to the data warehouse nightly, and the 3 main ways that the data is accessed by users is shown in Figure 1.

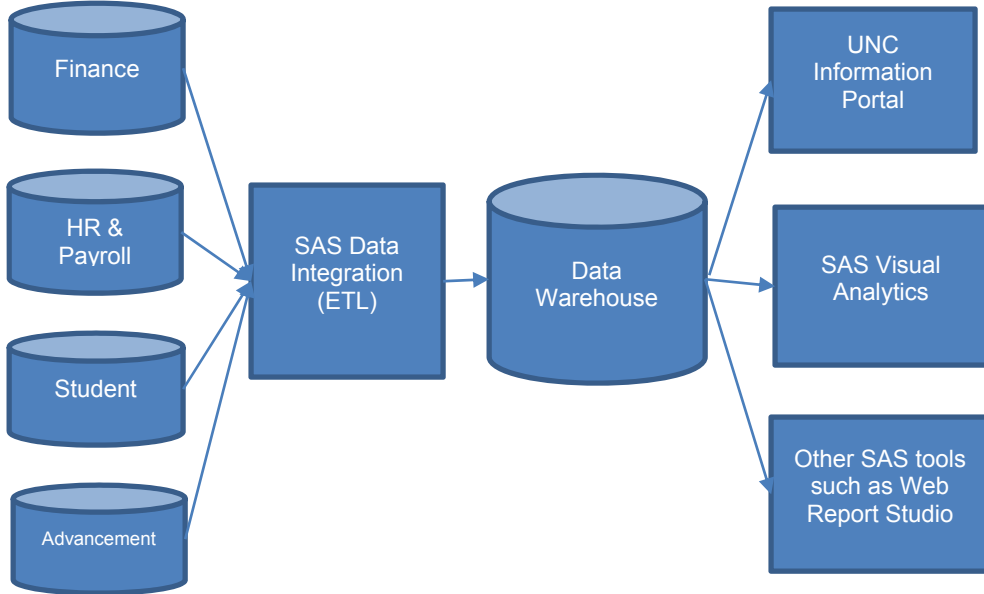


Figure 1. Data and Tools at UNC

PROFILE: REPORT CONSUMERS, REPORT WRITERS, DATA LOADERS

The user base estimates for full use of the system shown in Figure 2 include 5,500 report consumers, who will view one or more reports. In addition, roughly 200 users will be given additional capabilities to write their own reports, whether those are ad-hoc reports or formal reports to be shared to a particular business unit or the entire enterprise. Finally, we anticipate that we will grow to 100 or more users who are enabled to load their own data, and potentially mix that data with the central reporting data and reports. It is important to know that they are not simultaneous users, and not all of them will be heavy users rather we anticipate 50-100 simultaneous users. On the first day of go-live we had over 500 report consumers who came in to kick the tires between the hours of 7AM and 7PM, with a few stragglers coming through 1AM the next day. This and other metrics continue to inform us as to our critical hours, which appear to be 7AM-7PM.

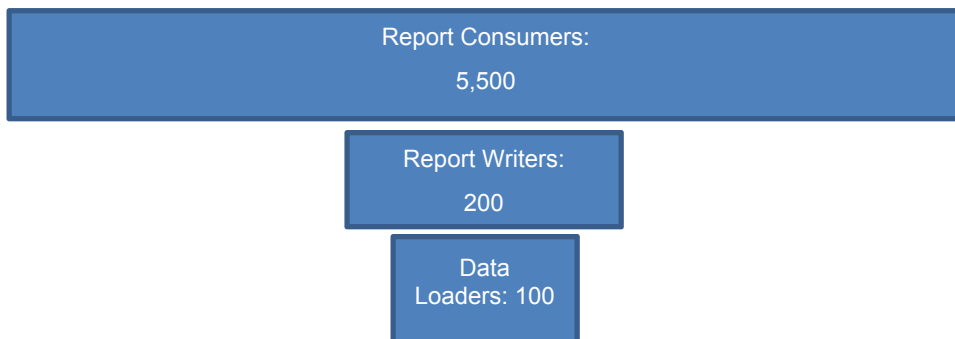



Figure 2. Report Consumers, Writers, and Data Loaders

ENVIRONMENTS: WINDOWS SERVER 2012R2 FOR EBI AND RHEL LINUX 6.6 FOR VA – DEMO, DEV, TEST, PROD

Our environment strategy is multi-pronged both in the choice of platforms and the division into separate environments for different audiences as shown in Figure 3



Data Loading and Central Report Development are done in Test and migrated to Production.

Ad-Hoc reporting can be done directly in Production.

Figure 3. SAS Environments at UNC

We had chosen the Windows Server 2012R2 platform for the SAS Business Intelligence and Data Management Advanced deployment. Some anticipated benefits for using Windows was the ability to run SAS desktop applications on the server to aid in resolving issues and for the larger number of support people at the university with Windows platform experience. Since Linux was a requirement for the distributed license of SAS Visual Analytics, UNC chose RedHat Enterprise Linux because of UNC technical experience with running enterprise software on Linux. Tradeoffs exist in running separate environments and different platforms. After experiencing separate environments for ETL and reporting, the current path is to co-join the two (Linux and Windows) in the same metadata repository and license to allow for cross-platform interactions (such as loading ETL data to Hadoop or directly to LASR from SAS DI Studio). Further alterations are being considering as the project moves forward. We hope to keep you informed with future updates.

The broad reach of this undertaking drove the requirement of a traditional Development -> Test -> Production lifecycle. The ETL code, supported by SAS Data Integration Studio was used inside of the SAS BI and DMA software. To minimize the impact on the end user, this code was to be developed and tested prior to releasing it to Production. As the administration moved forward, a need to test hot fixes and upgrades without impacting the developers came to light – thus, the need for a temporary copy of the Dev system was established to perform this work. This environment would be named Demo.

In SAS Visual Analytics, we have a need for a lifecycle of report development and promotion, but more importantly for data source validation before loading to Production. A number of times during development, accidental Cartesian joins of input data overwhelmed the VA system. Having the Dev -> Test -> Prod cycle in place aided in resolving these problems before promoting to Production and potentially impacting our customer. To get the data into memory for Visual Analytics in Production, the preference was to schedule data loads to the Hadoop file system then load to SAS LASR server. The development lifecycle included central office report development in the Test environment and then the promotion of the reports to production. This promotion was also done with respect to any any data loading, typically accomplished with the SAS Visual Analytics Visual Data Builder in Test promoted to Production. Currently, to meet ad-hoc reporting needs, trained users are allowed to create reports directly in Production.

To care for and feed these various environments and platforms, IT infrastructure and security concerns must be measured appropriately. All environments were placed in a private VLAN with 127 addresses, setup external routing through a proxy, protected the network with firewalls, and routed user traffic through an F5 BigIP network device.

INTEGRATIONS:

FIREWALLS AND NETWORK DEVICES

Business Case

In order to provide the necessary protection of the resources and data, firewalls should be employed to separate the public network, the campus network, the servers, and the databases. The end-users should access the software through a 'DMZ' network zone that provides a buffer and limited connectivity to the users. In turn the 'DMZ' network zone should have only the necessary connectivity to the servers and databases held in further network zones that we refer to as "Trust".

Caveats

Implementing complex and comprehensive firewall rules can be time consuming. The many SAS solutions use a fair number of ports, and getting those right and tested through the firewalls can occasionally cause services to become unavailable as work is being done throughout the organization on related servers, databases, or the firewalls themselves.

Technical Implementation

We place firewalls between the various network zones, and selectively allow traffic to pass from one zone to another. Figure 4 shows an illustration of how our users and their computers on our campus network are allowed only ports 80 and 443 into our "DMZ" to access our F5 BigIP Network Appliance, which takes the user requests and passes them to the servers on the other end, through another firewall into our SAS Firewall zone on ports like 7980 for web traffic, or port 8561 for metadata connections.




Figure 4. Firewall and Network Devices

SHIBBOLETH SINGLE SIGNON (SSO/SAML)

Business Case

Single Sign On (via Shibboleth SSO Provider) allows our users to reduce the number of times that they enter their password, offering more seamless integration between all systems behind SSO. It also offers a standard look and feel as well as a single URL that asks for credentials. This offers a higher security model, since our users are less likely to enter credentials in sites that are not SSO.

Caveats

The SSO login only lasts for 30 minutes to allow logon to new systems. Once logged in, systems will keep the user's session alive for potentially longer timeframes.

Technical Implementation

Our Single Sign On is only used for web applications on our services. Other authentication methods, for example with SAS Desktop tools, continue to require a user to provide a UserID and password. The implementation of the Shibboleth Single Sign On within SAS 9.4 and the web server is slightly different, and getting easier between versions. More details on how to configure this can be found by request from SAS Technical Support and also in the session SAS1385 - Federated Security Domains with SAS® and SAML at SAS Global 2015.

To get an idea of how this fits into the enterprise, the flow of a logon on the SAS Web Server and the integration with the SSO is shown in Figure 5. The steps are:

1. The User Web Browser contacts the SAS Web Server.
2. The Web Server redirects to the SASLogon web application, and redirects through the Shibboleth Service Provider sending the User web browser to the Shibboleth Single Sign-on provider.
3. The User Web Browser shows the user a logon page that is the same page as all other SSO enabled apps. The user enters their UserID and password.
4. The User Web Browser is redirected back to the SAS Web Server.
5. The User Web Browser is now able to interact with the SAS Web Server applications.

Once a user session has been established, SAS Token authentication is used to run jobs under a common UserID. This has advantages that allow central granting of permissions to this shared ID, but at the same time the disadvantage that there is no granularity of permission at the operating system level between users – the access must be limited through SAS Metadata permissions.




Figure 5. Flow of Single-Sign-On

MS ACTIVE DIRECTORY AND OPEN LDAP

Business Case

We want to be able to allow our users to use their central single user-id for logon to the SAS software to take advantage of all of the management and security benefits. To be successful with that, we also want to use the groups that are delivered in the central services to better scale user administration and security across all of the business units.

Caveats

We have both an Open LDAP and a Microsoft Active Directory here, and the Open LDAP data feeds to the Microsoft service. Our infrastructure integrates to both of these services, though there is a tendency for Linux to integrate with the Open LDAP while the Microsoft Servers and Desktops incorporate Active Directory. Since there is a dependency, and since we are on both platforms, we must strategically choose which SAS platforms consume each, and avoid confusion.

Technical Implementation

We started out by creating a directory structure in MS Active Directory and using that structure from both Windows and Linux. While the Windows Server/Desktop part is straight-forward, as most of those things integrate with MS Active Directory, on Linux we enabled the PAM integration above and beyond our typical OpenLDAP integration. We primarily use the PAM implementation for the consumption of groups from Active Directory.

An extract of some folders, user groups, and servers is shown in Figure 6. The example shows how we organized into several generic groups, with subgroups provided by each organizational unit included into the SAS Enterprise Citrix Users group, and each organizational unit designating one or more personnel to be the SAS Enterprise Departmental Admin, and a few select individuals that have attended appropriate SAS training in the SAS Enterprise Server Admins. We organized between production and non-production servers, without additional granularity between the Demo, Development, and Test environments in order to enable more access in the Non-Production environments to a broader base of users. Our Citrix group is held separately, with plenty of user access inside the Citrix groups, but overall access to all SAS users.

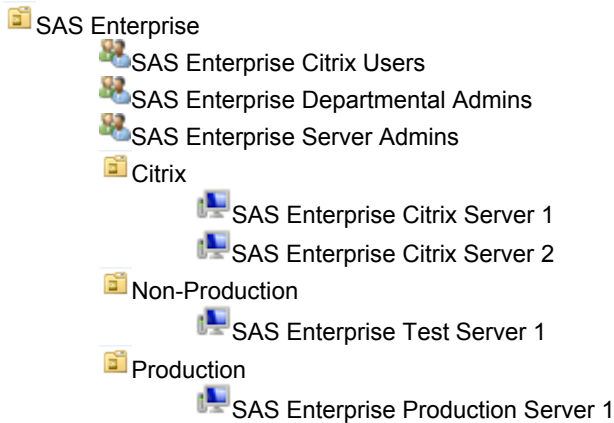


Figure 6. Example of Active Directory

To meet the specific additional requirements of the SAS Visual Analytics Visual Data Builder, each must have an operating system account that provides password-less SSH between all nodes in the cluster. In order to provide this, but limit the other types of access to the servers, a “lite” home directory, specific to SAS user profiles is created as each user logs in to SAS Visual Analytics for the first time. Automation also provides the password-less SSH between the servers, and any additional setup for the individual user. This reduces the administrative workload by removing the manual steps needed to accomplish these tasks. Additional automation is being added to a base script that imports users from the portal application discussed below, and is planned to include further synchronization with Active Directory and OpenLDAP.

UNC INFORMATION PORTAL

Business Case

An existing central portal for information should be leveraged to accommodate the new data warehouse data, and extended with the SAS reporting tools. Users that already have accounts in the portal should be granted appropriate access to view reports through the SAS tools such as Visual Analytics and Web Report Studio.

Caveats

Tradeoffs of this approach include the need to initiate a SAS session within 30 minutes of the logon to the Single Sign-On provider, otherwise additional navigation happens that breaks the Iframe. The ribbon style toolbar at the top consumes some vertical space on the webpage, and scrolling issues can occur because of the Iframe. While this provides the user with easy navigation, it does reduce usable space and eliminate the possibility of viewing multiple reports at once, which is already awkward within the SAS Visual Analytics Flash based user interface.

Technical Implementation

The web browser page is structured as shown in Figure 7, using a single Iframe to contain the SAS Visual Analytics Report Viewer. This is convenient for the user in that they are used to seeing the toolbar at the top and the reports below it. Several Icons/Tabs in the toolbar launch a SAS Visual Analytics Viewer session in the Iframe. Some launch to a menu of reports that in itself is a SAS Visual Analytics report, others have a menu of reports that take the user to a specific report inside of SAS Visual Analytics. One limitation is that the user may not readily view more than one report at a time, and this is the same limitation that is present for other elements in the portal.

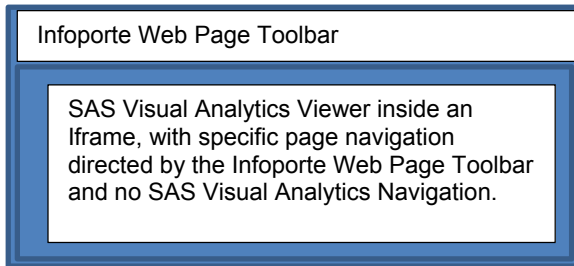


Figure 7. Conceptual Diagram of Portal Integration

In order to grant access to the SAS Visual Analytics Viewer, a synchronization process must occur whereby users of the Infoporte application are synchronized into the SAS Metadata. The source for this is slightly different than the typical Microsoft Active Directory synchronization, in that an Oracle table is used for the source information and the users are all placed into a single Users group in SAS. They may get additional group information applied manually for other metadata folders and assets that they will access outside of Infoporte. This process runs overnight as new users are given access to Infoporte. We have not yet implemented an Active Directory synchronization of users and groups, so this integration is helpful in giving user the necessary access. For the hundreds of users that will be doing report writing and data loading, we believe that we can manage the current demand manually for now, and then if necessary implement more Active Directory synchronization. Since the synchronization requires extensive SAS coding for a number of complexities, we are not able to do that easily out-of-the-box with the SAS software.

CITRIX FOR REMOTE ACCESS AND DESKTOP TOOLS

Business Case

In order to both contain data as well as to simplify the deployment and upgrade of desktop tools, Citrix is used as a solution at UNC. This minimizes the many issues that can occur with distributing desktop software to a large user base, as well as the various issues and incompatibilities of different versions of Windows operating systems. Furthermore, as the proliferation of Macintosh and Linux computers, as well as the trend of mobile/bring your own device continues, some valid business cases must be served through a server based resource to accommodate these varied user devices.

Caveats

Citrix is a complex and sometimes temperamental product that requires significant care and feeding. Since we have multiple resources working on this effort, and use it for other essential activities such as the development and operation of many critical IT assets, using it for SAS is simply an extension of those other capabilities. Of the SAS Windows Desktop products, so far we have only found that SAS Enterprise Guide does not play well on our small farm of Citrix servers. Because Citrix is server based, it is very important to ensure that the licensing from SAS includes licenses for both Windows Desktop operating systems as well as Windows Server operating systems.

Technical Implementation

We initially implemented two Citrix servers to host SAS Desktop Clients, and provide those to our users through our NetScaler network appliance as shown in Figure 8, ensuring strong security. These are both inside of the SAS network zone, so no firewalls openings are required to connect with the SAS servers. Two Citrix servers were chosen in order to provide redundancy. Since they are the only virtual machines, they are already highly available and redundant, so another reason was identified for the split. The need to upgrade the software and patch the machines from time to time, it was decided to use one as the "current" version and another as the "future" version. As the servers are upgraded, the choice of whether to use the current or future version of the desktop tools will be given to the user and hopefully automated in some way in the near term. The installation of the SAS software is simply to install the client software for the licensed products using the SAS Deployment Wizard. Other tools such as SSH tools, password

managers, and web browsers are also installed for use by the administrators and developers to access the necessary tools without requiring extensive firewall work and risk.




Figure 8. SAS and Citrix at UNC

A FEW OTHER INTEGRATIONS OF NOTE

While the above amount to the essential and highest value integrations when implementing an enterprise wide reporting and analytics solution with SAS software, these additional items are notable in that they all help just a little more and together represent additional value to the organization in time saved and more robust operation.

Network Attached Storage

The need to allow our users to get to any of their data and include it inside of the SAS solutions led us to adopt a strategy that we would allow users to attach to their NetApp server storage from the servers. This allows them to access any desktop data set, including CSV, Excel, Text, and SAS datasets, simply by providing the path to the network storage share starting with '\\server\volume\path'.

CA Workload Automation Scheduling

Since we control much of our overnight and weekend processing with our Computer Associates workload automation scheduling software, it was logical to install agents on our SAS servers and allow specific SAS jobs to be triggered in parallel and/or serially on the SAS servers. For example, the Finance jobs in PeopleSoft complete, then a trigger for the SAS ETL, followed by a trigger to load the resulting data into the SAS Visual Analytics server.

Network Security Scanning

We identified that our policy of scanning the network ports with abandon would result in stopped LASR servers on our SAS Visual Analytics environment. SAS Technical Support confirmed that one of the components was noted to be sensitive to this type of probing, and we worked around it by blocking the specific port ranges that were causing this issue.

Oracle, MySQL, MS-SQL, Postgres

Our strategy, in order to connect to many different types of databases, was to install the clients and develop connection strings that the various Library owners could use to connect to their database. This would avoid an administrator needing to define specific connections each time a user needed a new database. The one caveat is on Linux we found that the ODBC connections can only be defined by an administrator inside the ODBC configuration files. In addition, shared credentials are stored in Metadata for use by automated processes and authorized users accessing these databases.

Enterprise Virus Scan and Updates

We have an enterprise virus scan solution that ensures that our servers and user machines are as clear of malware as the products can detect. In order to run SAS, we had to put in file type exclusions as documented by SAS, and were able to share this out to the Windows servers through group policy.

We also use Windows Group Policy and Microsoft System Center Configuration Manager to auto-patch the Windows servers during non-use times that are scheduled twice a month – the first for Non-

Production, the second, two weeks later, for Production machines. Our Linux systems use a similar schedule and utilize Puppet software to accomplish the same goals.

Proxy Outbound Traffic

Part of our security strategy is to put the servers not only behind firewalls, but also on private network address space. Because of this and our lack of Network Address Translation (NAT) in this network, we made use of an enterprise level proxy server to allow access to resources outside of the university. The configuration is quite easy in Windows, but some programs, many in Linux, do not necessarily work with a proxy like the one we had, and we had to find alternatives such as when we experienced issues with the FTP for SAS updates, we switched to sharing the server files on a network share with Windows admin workstations that use WinSCP to overcome both the proxy issue on Linux as well as passive FTP connections through our protected network.

Backup Strategy

We use provided automation to write backup information from each server to a NetApp storage share using SAS Deployment Backup and Recovery Tool as a means to ensure backup. Our Tivoli Storage Manager is overkill (and quite expensive) for our server configuration since we are not storing any permanent data sets on the servers. In addition, we have a configuration directory for each of the servers that contains the necessary customizations in order to restore service in a disaster – which would have to be significant enough to interrupt service to our fully mirrored storage.

SAS Reports on Itself

We are using SAS to report on itself. For example, to provide an ETL Status and Completion Date/Time that is written by each job as it runs, and then the summary table is reported on to provide the latest status and also can be used for trending over time. The same holds true for our SAS Visual Analytics LASR loading, and also user history of accesses to the environment, and with our upgrade to SAS Visual Analytics 7.1, we will begin to see which reports are most popular and derive other metrics about the environment.

Enterprise Logging

We use an installation of Splunk for logging from our servers to assist in investigation and resolution of issues. The SAS servers fit into the model and all of the typical server data is logged to Splunk. We have observed that there are potentially hundreds of different types of log files, and further identification and inclusion of those SAS log files into our Splunk infrastructure is an ongoing effort.

SharePoint for Team Sharing

The last tidbit here is that we use Microsoft SharePoint as a tool for the campus SAS community and to store much of the documentation that we create about the system.

CONCLUSION

While we have spent a lot of time implementing the above integrations, we are still refining some and discovering others. Our strategy included buying all like hardware and shifting servers from areas of low utilization to areas that need more hardware – which we are currently doing in conjunction with SAS software upgrades – a highly choreographed effort that will involve at least 25 of our servers. We are also revisiting our row-level security implementation in Visual Analytics now that we understand better the strengths and limitations of this new technology.

The experience and lessons learned continue to guide and direct the path forward. Some advice we encourage you to consider when implementing an enterprise wide analytics and reporting solution would be:

- Do your homework – The more you know about your desired business process, data, and strategy, the better answers you will have to drive your solution to success. Ask others who have gone before you.

- Go through an Enterprise Excellence Center sizing exercise – Get to know the entire investment as you move forward. SAS software is powerful and can require equally powerful resources.
- Ask for help – SAS and their Technical Support have a wealth of knowledge and resources to answer questions and get experienced perspectives and solutions. SAS can also put you in touch with other customers who may have gone through similar business needs and how their goals were achieved.
- When in doubt, reach out – SAS enjoys partnering with their customers in solving their business problems. Even if you need help knowing what your business problem is, SAS can be a resource to help.
- Experiment – There are times when you need to think “outside the box” in order to get to the next level. Keep your assumptions in check and remain flexible with a plan that will undoubtedly change.
- Measure, find areas to mend... And keep moving forward – Always look for opportunities to improve. Marching forward is going to happen no matter what... look for efficiencies you can gain in adjusting project code, processes, and with experience and hindsight on your side, you can greatly improve your implementation and find new solutions.

Hopefully this is a decent roadmap to get you started. The hope is to share more of the experience as UNC further implements their SAS Solution.

For future installments of this paper, UNC continues to evolve the utilization and use of SAS for their ETL and reporting needs. The next edition would be to discuss outcomes of the lessons learned and how changing the architecture has impacted performance and efficiencies.

ACKNOWLEDGMENTS

There were so many people at UNC, SAS, and our various contractors who have contributed a great deal to the success of this project and the information contained in this paper. We wish to thank all of them for their efforts, because a project of this scope and size would not be possible without the efforts of the hundreds of people involved in this project.

CONTACT INFORMATION

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

Jonathan Pletzke
University of North Carolina at Chapel Hill
Pletzke@unc.edu

David Franklin
SAS, Education Practice
David.Franklin@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.