Paper 1425-2017

Migrating Large, Complex SAS® Environments: In-Place versus New Build

Chris James, UnitedHealth Group

ABSTRACT

SAS® migrations are the number one reason why SAS® architects and administrators are fired. Even though this bold statement is not universally true, it has been at the epicenter of many management and technical discussions at UnitedHealth Group. The competing business forces between the desire to innovate and to provide platform stability drive difficult discussions between business leaders and IT partners that tend to result in a frustrated user-base, flustered IT professionals, and a stale SAS® environment. Migrations are the antagonist of any IT professional because of the disruption, long hours, and stress that typically ensues. This paper addresses the lessons learned from a SAS® migration from the first maintenance release of SAS® 9.4 to the third maintenance release of SAS® 9.4 on a technically sophisticated enterprise SAS® platform including clustered metadata servers, clustered middle-tier, Secure Sockets Layer (SSL), an IBM® Platform Load Sharing Facility (LSF) grid, and SAS® Visual Analytics.

INTRODUCTION

SAS® platform migrations pose an interesting dichotomy to stakeholders. How can SAS® platforms keep up with the demands of tomorrow without jeopardizing the stability of today? What does an ecosystem look like that can satisfy both? This particular challenge is at the forefront of nearly all migration decisions. Addressing this challenge is difficult and typically causes disruption to the business, but it does not have to cause your organization to stay with the status quo. A successful migration is addressed through careful planning and a little help from SAS® Deployment Wizard (SDW) and SAS® Deployment Manager (SDM). This paper will discuss the lessons learned from an in-place upgrade and a new build with a metadata migration.

SETTING THE STAGE

UnitedHealth Group (UHG) has approximately 4,000+ internal SAS® users that provide solutions for nearly every business within UHG. The platform described in this paper is known as the SAS® High Performance Analytics (HPA) environment. The SAS® HPA serves 1,850+ internal users with a sophisticated topology distributed across 22 Red Hat Linux machines. The workload of this system is made up traditional batch and interactive jobs, plus hosts a full stack of SAS® Business Intelligence and SAS® Visual Analytics products. The workload management software that holds all the traditional SAS® jobs together is IBM®'s Load Sharing Facility (LSF). Lastly, this platform rarely has scheduled downtime and is intended to be available 24/7, 365 days a year. The purchase of new products and key enhancements to the third maintenance release of SAS® 9.4 drove the decision to migrate from the first maintenance release of SAS® 9.4. Display 1 provides a visual representation of the basic topology of the SAS® HPA.

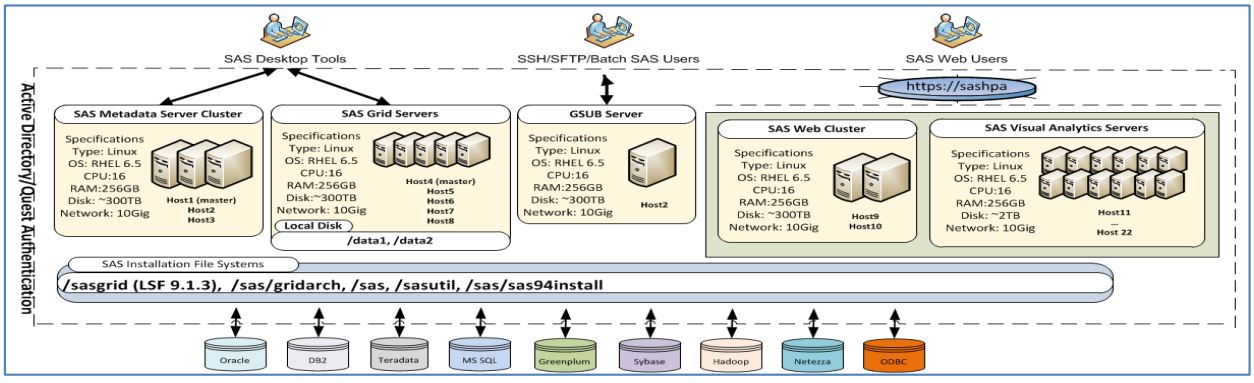

Display 1. Topology of the SAS® HPA

Two different elements impacted both the in-place upgrade and the new build and metadata migration. These two groups of products were IBM® Platform Load Sharing Facility and Process Manager and SAS® Deployment Wizard and Deployment Manager. Let me begin by introducing you to these two critical groups of products and explain how they impacted our upgrade.

IBM® PLATFORM LOAD SHARING FACILITY AND PROCESS MANAGER

IBM®'s Platform Load Sharing Facility (LSF) is used to balance SAS® jobs across heterogeneous infrastructure. LSF has a sophisticated framework that distributes jobs to the most available host based on the load and policies created by site administrators. The SAS® HPA relies on the power of LSF to manage a large, disparate workload for a large user base. A majority of the predictable workload is scheduled through IBM® Platform Process Manager (PPM) using a variety of time based triggers and file events. The remaining workload is attributed to jobs executed from SAS® Enterprise Guide, SAS® Enterprise Miner, and SAS® Stored Processes. The first step when moving from the first maintenance of SAS® 9.4 was to upgrade LSF to 9.1.3 to meet the minimum requirements for the products installed in the third maintenance release. This particular step was a predecessor to doing an in-place maintenance upgrade with SDW/SDM or choosing the alternative method of installing side-by-side with a new build. Because of the dependency to be at LSF 9.1.3, a change control was scheduled and the process to upgrade LSF began.

Upgrading LSF from 8.2 to 9.1.3 was actually pretty straightforward. The entire LSF cluster, including thousands of scheduled PPM jobs, are upgrade to 9.1.3 with a few simple steps. The good news is that the key configuration items stay the same between the two versions. The install.config file used during the original LSF 8.2 install can be reused to populate the 9.1.3 version of the file. Output 1 lists the installation command that must be issued by root to install LSF 9.1.3.

```
lsf9.1.3_lsfinstall/lsfinstall –f install.config
```

Output 1. LSF Installation Command

The LSF upgrade was validated and the changes to the platform were accepted by the business. The first problem of many to come would appear during the weeks following the LSF upgrade. Some scheduled jobs starting missing their scheduled runtimes completely, some scheduled jobs would spawn multiple instances of the same job, and some scheduled jobs would simply hang without ever receiving a job id from LSF. These problems were escalated and examined by both SAS® Institute and IBM. The PPM option JS_FILEAGENT_SENSITIVITY was changed from 30 to 60 to correct the duplicate job instances.

One suggestion that did not resolve any issues was to move our entire LSF installation and configuration from EMC Isilon to NetApp because of Isilon's limitation on handling file locking over nfs3 in a multi-machine topology. This move not only did not resolve any issues, it actually caused more problems in our environment because NetApp is internally configured to use Lightweight Directory Access Protocol (LDAP) which meant that all user ID and UID number differences between LDAP and Active Directory (AD) would prevent a user from accessing the LSF profile unless an override was added to the NetApp usermap.cfg or the user id and UID number were synchronized to be the same in AD. This was significant because over ten percent of the users have mismatched user IDs between LDAP and AD. Fortunately, SAS® programming and LDAP queries were used to find all the users with mismatched IDs and a list was provided to the enterprise NAS team for quick resolution. This did, however, pose an ongoing support problem with new users that came onto the platform with mismatched IDs. After several months of countless help desk tickets and no progress on issue resolution, the LSF installation and configuration was moved back to EMC Isilon.

Part of the migration strategy to upgrade LSF from 8.2 to 9.1.3 was to move the LSF master host responsibility from a host that executes several key functions, including jfd (PPM), PostgreSQL database, SAS® Object Spawners, and normal grid workload, to a different grid node. The separation of jfd and the LSF master host caused some jobs to miss their scheduled runtime. This separation may not be an

issue for some SAS® platforms, but it certainly was an issue for the SAS® HPA. The resolution for this issue was simple; move the LSF master host back to the original.

The last LSF issue that plagued this system for months and months beyond the maintenance release upgrade for SAS® 9.4 was related to PPM jobs hanging without receiving a job id from LSF. More specifically, PPM would start a job based on a scheduled time or a trigger file. The communication between PPM and the lsb.events in LSF was not consistently occurring, which resulted in PPM never receiving a notice that the LSF job started, even though it did in most cases. This issue would also cause dependent flows in PPM to wait because the notice was not received that the LSF job started and ended. Our team was directed to move the LSF installation and configuration directory to a file system that could handle locking and could keep up with the constant IO necessary to keep up with all the transactions that occur in lsb.events. This ended up being a red herring because moving the LSF installation and configuration directory to another file system had no impact on the outcome of this issue. The resolution for this issue occurred when the operating system was upgraded from Red Hat 6.2 to 6.5. This upgrade immediately resolved the issue without any additional tuning to LSF or PPM.

```
1.) Update JS_FILEAGENT_SENSITIVITY from 30 to 60 to correct the duplicate
    job instances.

2.) Keep the LSF master host and jfd (PPM) on the same physical host.

3.) Upgrade to Red Hat 6.4 or higher to resolve the communication issues
    between PPM and LSF's lsb.events.
```

Output 2. Key Resolutions to Issues with LSF 9.1.3

SAS® DEPLOYMENT WIZARD AND SAS® DEPLOYMENT MANAGER

The SAS® Deployment Wizard and SAS® Deployment Manager play a foundational role in any installation, configuration, and maintenance upgrade. SDW is intelligent enough to discover upgrades that are available in a SAS® Software Depot. This makes installing new products or upgrading existing products seamless. SDW also has a number of impressive options including recording response files, partial prompt and quiet installations, and automatically installing the latest hot fixes during the installation. SDM, on the other hand, is only as intelligent as your initial configuration. All the choices you select during the initial configuration of SAS® are propagated throughout the SAS® environment and subsequent executions of SDM will continue to use those initial selections, even to the detriment of the current SAS® environment. This concept is critical for anyone that has a complex SAS® environment with post configuration changes.

In the routine of designing the architecture, finding funding, and outlining how a SAS® system will intrinsically satisfy business needs, the approach that SAS® administrators will take to perform maintenance upgrades falls to the bottom of the priorities. In fact, I would dare to say that maintenance upgrades are thought of as common place, like installing a group of hot fixes. The key priorities that are routinely considered in a large, enterprise platform include: performance, security, availability, resilience, consistency, and user experience. These elements can be woven into the fabric of the SAS® ecosystem. The SAS® Metadata Server is easily clustered during the original configuration of SAS® and performs very well in a large enterprise setting. From a Base SAS® perspective, much of the goals can be achieved through sophisticated load balancing using IBM® Platform LSF or SAS® Grid Manager for Hadoop. The SAS® Middle Tier is what poses the most difficult challenge. SAS® provides a well-documented method of deploying the SAS® Middle Tier using advanced techniques to satisfy all the high-level goals of our organization. The recommended approach is to deploy a vertical or horizontal cluster to enhance the SAS® Middle Tier performance and reliability. A vertical cluster is the deployment of multiple web application servers on the same machine whereas a horizontal cluster is the deployment of multiple web application servers on multiple machines.

During the initial build of this platform, SAS® Professional Services and administrators from UHG took this process one step further by adding a horizontal cluster and making as many additional services highly available as possible. This included adding an additional SAS® Web Server and JMS Broker on the second member of the horizontal cluster. An F5 switch was also added to simplify the URL and provide a layer of branding to identify this SAS® environment. These post configuration changes would enhance the platform in the present, but would create an impasse for future maintenance upgrades. With all of this in mind, it is probably a good idea to lay out a visual representation of the current state and the desired future state of the SAS® Middle Tier in Display 2.

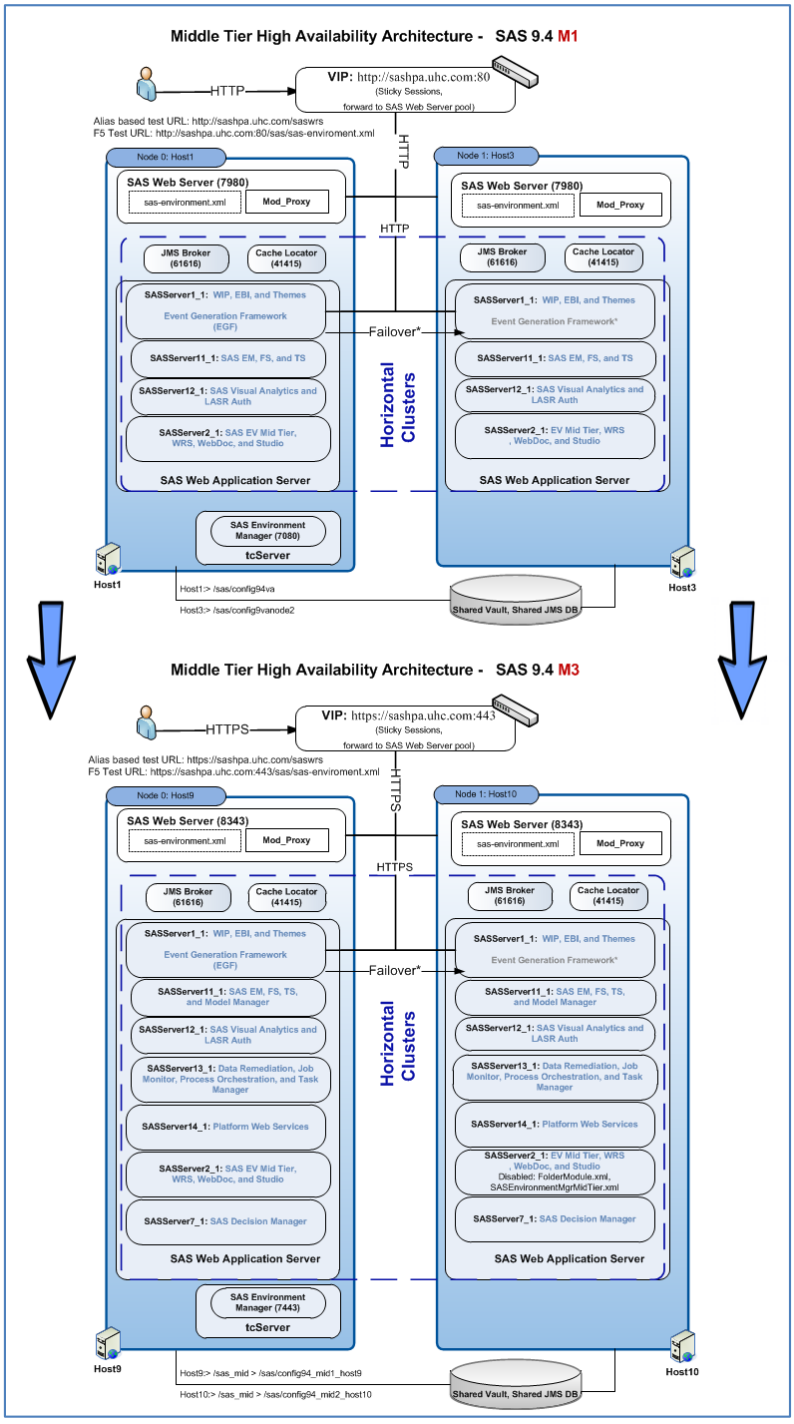

Display 2. Current and Future State of the SAS® Middle Tier

Display 2 also highlights an additional configuration element in the future state of the SAS® HPA to enhance web security, which is to SSL to the SAS® Web Server, SAS® Web Application Servers, and SAS® Environment Manager. The post configuration additions to the SAS® Middle Tier and the requirement to add SSL to the highly available topology would ultimately determine the outcome of the migration decision. This is discussed in detail later in this paper.

COMPARING THE MIGRATION OPTIONS

This paper will outline both options that were considered (and executed) to move the SAS® HPA from the first maintenance release of SAS® 9.4 to the third maintenance release. The decision between upgrading in-place and building a new platform is a substantial decision that impacts all stakeholders: customers, businesses partners, leaders, and SAS® administrators. In no way should the experiences at UHG serve as an industry standard or the only path for a successful migration. The migration strategy for UHG was the best option given the complexity and magnitude of the SAS® HPA. The next section will outline the key considerations, successes, and failures for both migration strategies.

IN-PLACE MIGRATION

The in-place upgrade is the fastest, least disruptive approach to staying current with maintenance releases from SAS. This approach avoids any form of an actual migration and delivers new enhancements quickly. The upgrade occurs simply by launching the new SAS® Software Depot for the third maintenance release of SAS® 9.4 and SDW and SDM do the rest. The prior is absolutely true for a number of SAS® environments, but not for the SAS® HPA. The complexity of the SAS® HPA, with products ranging from IBM® LSF, to SAS® Business Intelligence, SAS® Visual Analytics, metadata clustering, and middle tier high availability, made this upgrade a daunting task.

Although all options were considered, one path was more attractive to stakeholders than others. That option was to use the combination of SDW and SDM to do an in-place upgrade to the third maintenance of SAS® 9.4, which is the global standard for upgrading SAS® in most scenarios. This method of installing the maintenance release seemed straightforward as we had done many of these in the past. What could possible go wrong? As you could imagine, we faced many challenges with the in-place upgrade and we overcame all of them except for one.

Advantages

There are several advantages to doing an in-place upgrade from one maintenance release to another. The key advantage is time to market verses building a new system and having to migrate metadata content and file system data to the new system. This cannot be overstated as the SAS® HPA has thousands of users, jobs, metadata folders, projects, etc. There is also less risk introduced into the integrity of the original installation by choosing to upgrade what is already configured. All of the SAS/ACCESS Interface modules will continue to work as designed, plus SDM will preserve and update metadata content and keep custom settings on application server contexts. The other big advantage to upgrading in-place is the cost savings from having to buy new hardware and/or investing months of administrative hours into rebuilding the entire SAS® environment.

Key Consideration: Always Perform the Maintenance Upgrade on a Lower Environment

Performing the maintenance upgrade on a lower environment before upgrading production is always the global standard if a lower environment exists. The SAS® HPA disaster recovery (DR) environment was also at the first maintenance of SAS® 9.4 and was taken through the entire process without any issues. All the products that had updates available were successfully upgraded. New products that were either recently purchased or released were also installed and configured successfully, including SAS® Studio and SAS® Model Manager.

> *Note: The new requirement to add SSL to enhance the SAS® Middle Tier security was not considered during the initial maintenance upgrade. This requirement would be satisfied after production was updated to the third maintenance release of SAS® 9.4.*

The problem with the installation and configuration of the SAS® HPA DR environment was that it was not an exact match in terms of the product mix nor topology to the production system. The SAS® HPA DR environment did not have any type of SAS® Middle Tier clustering, horizontal or vertical. It obviously did not have an F5 switch incorporated either, nor did it have multiple SAS® Web Servers and JMS Brokers. The results from this test ended up being a false positive. Looking back at the decision to attempt an in-place upgrade, this was the biggest mistake that was made.

Where Everything Went Wrong

After months of planning and testing, it was time to perform the upgrade to the SAS® HPA over a 52-hour outage to the platform. The team was ready to begin the grueling outage that would last an entire weekend. The first step that was taken was to shut down all products on the platform and capture a full tar of the SAS® installation and configuration directories just in case something went wrong during the in-place upgrade. This ended up being the best decision made in the context of the planning and execution of the in-place upgrade. The installation of the new products continued as planned and was successful. All the trouble started during the configuration of the SAS® Middle Tier on Stage 8 when SDM tried to configure the SAS® Web Infrastructure Platform. The exact error message that was displayed is available in Display 3. Several attempts were made to perform the upgrade over the course of 52 hours and all of them resulted in this error message. SAS® Technical Support stood by us during the entire ordeal, but was not able to resolve this issue.

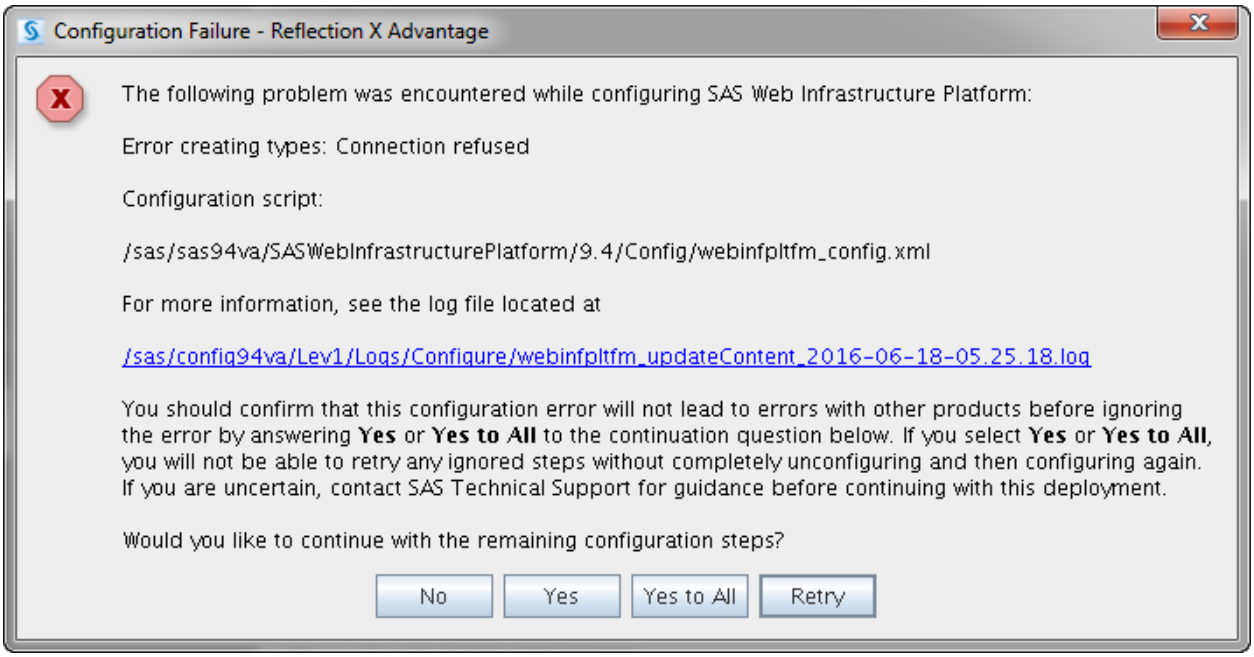

Display 3. Error Message During In-Place Upgrade

Explanation of the Issue

The issue boiled down to a misunderstanding between SDM and what was actually configured on the SAS® Middle Tier. SDM expected the SAS® Content Server, which is part of the SAS® Web Infrastructure Platform, to be at a specific hostname and port combination, but because of post configuration changes to make the SAS® Middle Tier highly available, the SASServer1_1 web application server was listening on the proxy address and port instead of the actual hostname and original port. That would explain why SDM stated "Error creating types: Connection Refused." At the time the issue was experienced in June of 2016, there was no known solution. SAS® Technical Support provided a script to try and accomplish the load into the SAS® Content Server, but nothing worked. Without drawing out all the nitty details, we simply stopped all products and restored the SAS® installation and configuration from the tar that was created at the beginning of the outage. As you would imagine, the next step would be to technically consider a new build of some type since the update in place was not successful.

NEW BUILD

With a failed in-place update in hand, frustrated customers, and management asking when the new version of SAS® would be available, we had to make a decision about building a new platform. There was not valid business justification to request capital to build a new 22 machine SAS® platform with the exact same topology of the current production system. In fact, there really wasn't a compelling reason to consider new hardware other than separating the production workload from the new installation of the third maintenance release of SAS® 9.4. Building a new SAS® HPA on the current hardware would certainly have some advantages over an in-place upgrade, but there are also additional challenges from pursuing this avenue. The next section will review the high-level advantages and the obstacles that would have to be overcome.

Advantages

The most obvious advantage to building a new SAS® platform instead of upgrading in-place is that this method spreads the amount of risk the business assumes over a period of time by giving the business users an opportunity to fully test the update to the third maintenance before the cutover. An in-place upgrade, with no development or test system available, would mean the customers would just land in the third maintenance without any prior testing. The in-place method of upgrading is a high risk, high reward approach to upgrading a platform of this kind. Another key advantage to a new build is that database drivers can be updated to modern releases and tested, which included MS SQL Server and Netezza for our platform. Lastly, the new build approach to migrations allows additional time to revisit the architectural choices that were made during the initial installation. In our specific deployment, there were Greenplum database servers that were not being used. These hosts were repurposed to support the new SAS® Middle Tier horizontal cluster because one of the original SAS® Middle Tier machines was undersized for the number of web users on the platform.

Obstacles to Overcome

A number of key obstacles exist to building any SAS® platform, but many of these are exaggerated because of the magnitude of this system. The wide-ranging menu of SAS® products that are supported on the HPA would mean months of work to mirror the current configuration. A key assumption of the entire project was that the new system would need to mirror all the functionality of the old system, which means all the security, custom configuration for multiple application servers, and the customer experience would have to be exactly the same as the first maintenance release of SAS® 9.4. The new build approach would mean a completely autonomous SAS® installation and configuration on the same general set of servers and primary storage device.

The initial stages of the planning phase included a full port review of the original installation, configured at the Lev1, compared to the new installation which would be configured at Lev4. Most of the potential port conflicts are handled by simply incrementing the level of the configuration. Other port conflicts were handled by using the HTTPS protocol instead of HTTP on the SAS® Web Server, SAS® Web Application Servers, and SAS® Environment Manager. The new port changes would mean a subtle change for the 1,850+ users accessing either the SAS® Metadata Server or the SAS® Middle Tier. One subtle change that would be worth its weight was to implement an SSL layer on all SAS® Middle Tier applications. New ports and the use of SSL certificates would also mean a change to the F5 switch that manages the highly available SAS® Middle Tier.

Another significant challenge of the new build approach on the same physical machines is the potential for the new installation and configuration to interfere with the production workload. This is especially true for the SAS® grid nodes, the SAS® Middle Tier, and SAS® LASR nodes. The SAS® grid nodes where performance tested using the same SAS® options in the sasv9_local.cfg, which included MEMSIZE, SORTSIZE, BUFNO, BUFSIZE, UBUFNO, UBUFSIZE, JREOPTIONS, and AUTOCALL libraries. The machine and disk settings would not need to be changed because they were already tuned for the production workload. Because of the sheer number of Java Virtual Machines (JVMs) in the SAS® Middle Tier, an architectural move was made to isolate the new SAS® Middle Tier on two Greenplum database servers (database shut down). These machines both have 256GB of RAM and are in a better position to support the large user base. The last concern we had with interference was related to SAS® LASR and

the TKGRID installation. TKGrid, TKGrid_REP, and TKTGDAT, were already installed and configured across twelve LASR servers. Adding the third maintenance release of SAS® 9.4 would mean adding TKGrid 3.3 on the exact same nodes, but using a different path. The Lev4 increment made during the new installation and configuration would take care of the ports, but we were not sure how the inter-communication between the head node and the worker nodes would work with two implementations of LASR running on the same set of servers. The good news is that the new side-by-side installation of SAS® LASR worked without issue. The only thing that needed to be monitored by having multiple versions of SAS® LASR on the same physical machines is the memory usage of the tables that are loaded into LASR. We certainly did not want to breach the 80% memory mark, which is the value set for the Data Loading (%) value on the SAS® Analytics Server.

Key Consideration: Slow Time to Market and Metadata Migration

The new build approach has a number of disadvantages, but the most critical two are the time to market and the need to fully maintain an exact replica of metadata. The time to market when doing an in-place upgrade using SDW/SDM is extremely fast. The time to move from one maintenance release to another would take about 52 hours. The time to market for a new build took our team about three months to fully install and configure SAS, add high availability with the help of SAS® Professional Services, customize the SAS® metadata settings, add our complicated security structure, migrate content, and test the maintenance release all while closely monitoring the current production system on the same servers. This method also requires collaboration with the different customer groups to make sure their applications will continue to function as designed in the new system.

The other key consideration was the process of migrating vast amounts of SAS® metadata and continuing to keep the metadata current until the cutover weekend. The plan was to use promotion packages created in SAS® Management Console to replicate the SAS® metadata during the cutover weekend. To allow the users to test the process, and for our team to practice the metadata migration, the decision was made to import the entire metadata structure several weeks before the cutover weekend. The metadata export from the current production SAS® HPA took about four hours to complete, but the process finished without issue. Importing the large metadata packages was not as successful. In fact, the metadata import into the new system corrupted the SAS® metadata server. The error that was received stated that the connection to the SAS® Metadata Server was disconnected. The log file for the SAS® Metadata Server slave host just abruptly ended. It is not until the SAS® Metadata Server is restarted that the corruption is found. The problem notes listed in Output 3 speak to several SAS® metadata corruption issues thought to be related to the issues we experienced, but it did not address corruption related to importing large packages.

```
1.) Problem Note 57428: In a SAS® Metadata Server cluster, the master node
    does not send redirection request to a slave node that is a part of the
    quorum

2.) Problem Note 59039: The SAS® import procedure causes an out of memory
    condition, corruption of metadata, or a loss or quorum for a clustered
    SAS® Metadata Server
```

Output 3. Hot Fixes to Resolve SAS® Metadata Corruption from Large Packages

Unable to import large amounts of metadata in one process, the only alternative was to import the SAS® metadata piece by piece. This process ended up taking several days to complete and entirely changed the strategy to import the full metadata again on the cutover weekend. The alternative approach was to give a deadline to production users, which was one week before the cutover weekend, to make metadata changes. The week of the cutover weekend was spent importing the SAS® metadata one last time. This process is not ideal, but it was the only method that would avoid metadata corruption and ensure the metadata on both systems would be in sync after the cutover weekend.

CONCLUSION

The decision to do an in-place upgrade verses building a new SAS® system is difficult and dependent on the circumstances of your SAS® environment. There is no golden rule, no industry standard, and no easy method. Each path is accompanied with a large amount of risk, but the risk should not make you stagnant. If your SAS® platform is good, it will eventually become outdated, slow, and behind the times if changes are not adopted. A key takeaway from this paper is to think about upgrades when you are building the SAS® platform the first time. If at all possible, there should be no post configuration changes that occur outside of the scope of SDM. If post configuration changes are absolutely required, these changes should be undone or unwound at the time of the maintenance upgrade to allow for a successful in-place migration. Depending on your environment, this can lead to extensive re-work and likely will extend the outage time to the customers.

REFERENCES

IBM Knowledge Center. 2017. “Introduction to Platform LSF.” Accessed February 22, 2017. https://www.ibm.com/support/knowledgecenter/en/SSETD4_9.1.3/lsf_foundations/lsf_introduction_to.html

SAS Institute Inc., SAS Deployment Wizard and SAS Deployment Manager 9.4: User’s Guide, Cary, NC: SAS Institute Inc., 2015.

SAS Institute Inc. 2015. SAS® 9.4 Intelligence Platform: Installation and Configuration Guide, Second Edition. Cary, NC: SAS Institute Inc.

SAS Institute Inc. 2016. SAS® 9.4 Intelligence Platform: Middle-Tier Administration Guide, Fourth Edition. Cary, NC: SAS Institute Inc.

ACKNOWLEDGMENTS

A special thanks goes to Diane Nieman for her effort in editing this paper. I would also like to thank all the members of the Optum SAS Team and the leaders in my organization for all the support and opportunity that exists at UHG.

CONTACT INFORMATION

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

Chris James
UnitedHealth Group
612-632-6918
Christopher.James@optum.com