

SAS® Admins Need a Dashboard, Too

Tricia Aanderud, And Data Inc.
Michelle Homes, Metacoda

ABSTRACT

Why would a SAS® administrator need a dashboard? With the evolution of SAS® 9, the SAS administrator's role has dramatically changed. Creating a dashboard on a SAS environment allows the SAS administrator an overview on the environment health, ensure resources are used as predicted, and provide a way to explore. SAS® Visual Analytics allows you to quickly explore, analyze, and visualize data – so why not bring the two concepts together.

In this session, you will learn tips for designing dashboards, loading what may seem like impossible data, and build visualizations that guide users toward the next level of analysis. Using the dashboard, SAS administrators will learn ways to determine the system health; and how to take advantage of external tools, such as the Metacoda® software, to find additional insights and explore problem areas.

INTRODUCTION

Once a SAS environment has been deployed, how are the user activity; the data accessed, metadata managed and report usage monitored? How can the business gain an insight into their investment and deliverables and how can a SAS administrator meet the business requirements while also gaining an insight into the system resources and performance. It is valuable to an organization to have an overview and to monitor the end user activity of a SAS Enterprise Business Intelligence (EBI) platform.

This paper has two sections. In the first section, we discuss the process for building a dashboard. To demonstrate the steps, the dashboard customer is a team of SAS administrators and associated business users. The data is collected from activity and logs within the SAS EBI environment and a third party tool, Metacoda Plug-ins. In the second section, you will learn how the dashboard process were applied to a department rollout and how to use the dashboard and supporting data to learn more about the system usage. We used SAS Visual Analytics for the dashboard creation and data explorations.

You could apply the concepts outlined in this paper to all SAS solutions using the metadata platform. The audience for this paper is SAS administrators and business users, but it may also be of interest to IT administrators, business analysts, project managers, team leaders and executives who are looking for a way to gain insights and forecasts in the activity and performance of their SAS environment.

SECTION ONE: DASHBOARD CREATION PROCESS

It may not be clear to some that a SAS administrator and SAS EBI users are different roles, with little in common, besides the tool. A natural disconnect further develops as each group forms an idea of what the other does without any actual learning, which might cause friction as they try to solve issues. One solution is to have the groups work together to create a dashboard based on the common tool, then each group is able to learn about the system from the other's stance.

By working together, the teams are able to monitor activity and resources and become more proactive in identifying issues quickly, recognizing patterns for potential end-user training and finding opportunities for resource maintenance, effectively reducing an organizations business risk.

UNDERSTANDING DASHBOARDS

As SAS administrators and SAS BI business users have different viewpoints, a good place to start is to ensure there is a shared understanding of what a dashboard is. When asked to design a dashboard, you must spend some time with the user understanding their needs and goals. Many users ask for a dashboard and use the term in a generic way to mean a web-based report with charts and graphs that can link to other web pages. A dashboard is more than a report because it involves conclusions.

In the *Effective Dashboard Design* book, La Grouw notes that you can differentiate a dashboard and report by their core attributes [ref 1]. A report does not lead users to a conclusion; it simply presents facts, which allows users to make a judgment. Consider the example in Figure 1, all the user knows is what the sales were for each month and sales improved by the end of year after taking a sharp dive in the spring. These are all useful facts but what we do not know from looking at these facts is whether it is a good thing or not.

A dashboard can provide context for facts; thus, providing that basis to form a conclusion. Effective dashboards can help a user to understand quickly what is happening and what needs to happen next. In Figure 1, consider the data in the right pane when presented with the sales goal for each month or even when presented as a total with the dashboard bullet chart. Presenting the facts, in this way, provide an immediate understanding of what is happening, and any actions that is required in response.

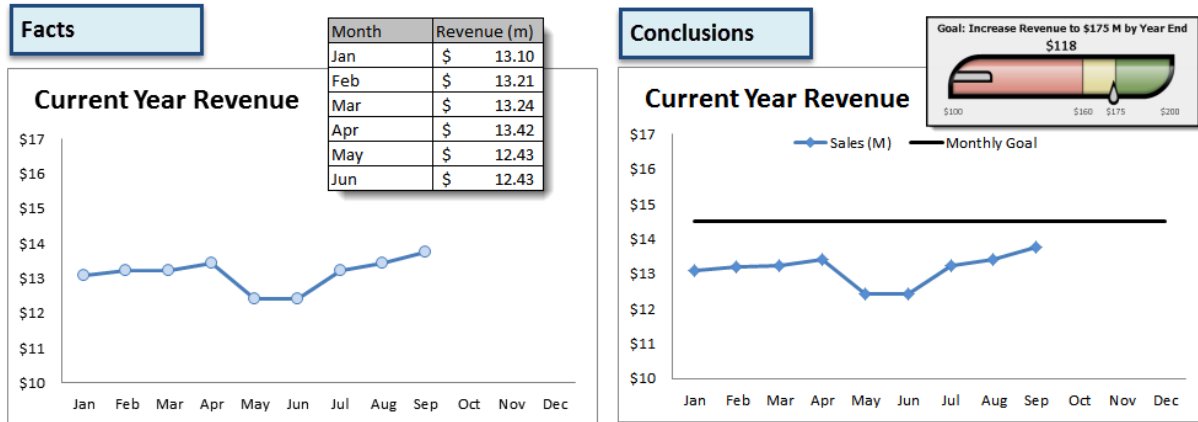


Figure 1 Understanding facts and developing conclusions

Dashboards help users understand if a process or project is running smoothly. Typically, metrics are centered on some performance aspect, such as strategic objectives, time lines, or quality result. Keep in mind, the information does not present the next level of detail, which is what or why events are happening.

La Grow suggests that dashboards be divided into two categories: Management and Operational. These categories are based on the information presented and the target audience.

Management dashboards are used by different levels of management and are based on the end goal. You can divide these dashboards into two broad categories: *strategic* and *tactical*. Strategic dashboards align with the organization's strategic objectives while tactical dashboards address more localized areas of concern. These dashboards use historical data rather than real time data for trend analysis.

Operational dashboards are near real time focused and aimed at front-line staff. In *Profiling Operational Dashboards*, Eckerman notes two operational dashboard types: *Detect and Respond* and *Incent and Motivate*. The huge wall of monitors that NASA staff use during a space shuttle mission is a good example of a Detect and Respond dashboard. Alternatively, call centers use Incent and Motivate dashboards, generally displayed on a wall, to encourage friendly competition or to allow managers know who needs assistance or training [ref 2].

DASHBOARD BUILDING PROCESS

The dashboard building process includes the following steps:

1. Determine what the user wants to measure.
2. Create sample layouts and build multiple prototypes toward the final design.
3. Collect the supporting data.
4. Create the final dashboard.
5. Observe the users interact with the dashboard to ensure adoption.

To illustrate, let us assume that we are building a tactical management dashboard for the SAS administrators and business users. We will explain how each step works to guide you through the process.

DETERMINING WHAT THE USER WANTS TO MEASURE

When starting a dashboard design, you might find it most useful to work with the users in a brainstorming session around a white board. This session might cover topics such as goals, pain points, and any existing metrics. Encourage the users to describe things that could cause issues or what success looks like. If you are talking to SAS administrators and business users, you might hear examples like the following:

- Can we identify memory usage on a server that is consistently swapping when accessing a report?
- How can we ensure that the system is being used and adopted?

- When is the best time to apply hot fixes?

In summary, users want to be proactive about the health of the environment and ensure that it can meet the organization’s goals.

How do you Measure a SAS Environment?

By the end of the brainstorming session, you notice patterns beginning to form and you can identify measurements to support the organization’s goals. You might see three main concerns: users, data, and system performance (shown in Figure 2).

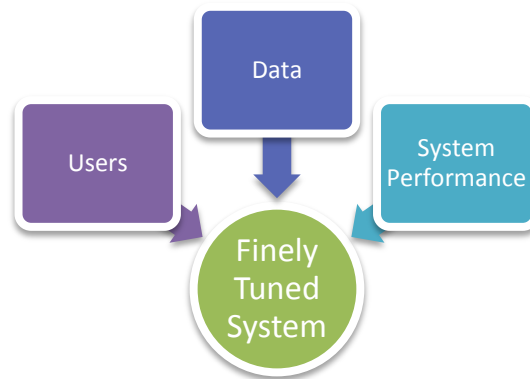


Figure 2 Concerns for System Administration

The following table lists three major areas of concern and suggested measurements. There may be overlap in these categories, so you will need to determine which goal is the best match your needs. At the end of the discovery process, you should have a list of goals that allow you to find the data sources.

User Management	Data Management	System Performance Management
<p>Relates to how the end-users interact with the environment. Examples:</p> <ul style="list-style-type: none"> • Ensure new department employees login within 5 days of training completion • Encourage user adoption by proactively contacting users who seem to have difficulty logging in • Identify users without any activity for 60 days • Ensure users have the appropriate access as per the security plan 	<p>The data aspect here relates to the databases or data files, reports and associated metadata. Examples:</p> <ul style="list-style-type: none"> • Remove unaccessed reports after 180 days of non-activity • Conduct database maintenance at non-active times • Archive the least used tables and folders • Monitor an inventory of metadata on folders, libraries, reports, stored processes, tables, and servers 	<p>Relates to managing and monitoring hardware and software in relation to the SAS environment. Examples:</p> <ul style="list-style-type: none"> • Ensure that the hardware can handle user activities • Limit stored processes to 1 minute of elapsed time • Determine the optimal time to apply hotfixes • Monitor system resources such as CPU, memory, and input/output

CREATING SAMPLE LAYOUTS

As the goal process continues, you can start thinking of the dashboard layout. Often dashboard designers find the layout of the indicators challenging. Using a tool like Microsoft Visio, you can play with indicator layout to consider how to group and display the data. Then you can present these ideas to the team for discussion.

Often when users see the indicators together, they start to realize what may not make sense. Some goals do not really belong on the dashboard. For instance, the administrators want to perform maintenance activities during non-active times. It is not really clear what goal we could use on the dashboard or how we position the measurement to add a conclusion. Perhaps the user activity really is information that belongs in a web report that allows the user to explore at will. In this case, the data is gathered but presented in a different report section.

All of the indicators from the SAS Visual Analytics Designer are box-shaped even if it is a round pie chart. In Figure 3 you can see a sample layout with suggested exploration paths. Remember the “01 area” is where people look first – so your most important indicator goes there. For layouts, you generally want the most important indicators larger while the lesser indicators are smaller. As you review Figure 3, you can see what elements are considered the most important, how indicators are linked and where the user would go for the supporting information.

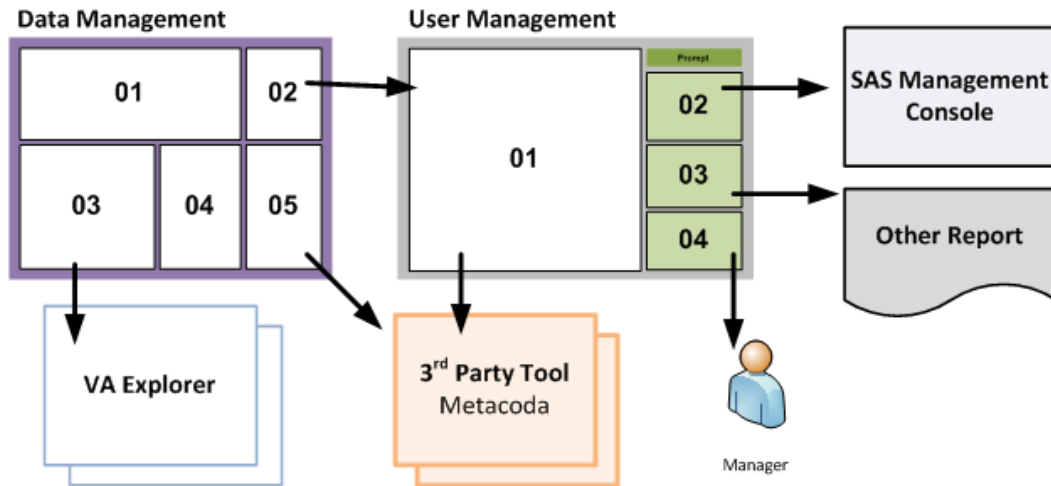


Figure 3 Dashboard layouts and information paths

COLLECTING THE SUPPORTING DATA

After determining the goals for the Admin dashboard, the data planning, collection and processing step begins. You should not discount goals where the data is hard to find. Instead, determine if the goal is stated correctly. It could be that the user needs to redefine the goal to better support the data. It's also possible that your current environment doesn't collect the information and you might want to consider alternate sources for the data, such as manually collecting data with a spreadsheet or a third party tool.

Where does the data come from in a SAS Enterprise Business Intelligence (EBI) environment?

A SAS 9.4 EBI architecture typically consists of 4 tiers: data tier, server tier, middle-tier and client tier. All of these tiers work in harmony with the SAS Metadata Server being the central heart of the environment, storing and saving metadata and responding to requests throughout the platform.

When a SAS client makes a request, such as running a SAS program in SAS® Enterprise Guide®, clicking on a SAS Web Report Studio report or running a stored process, the action, the SAS response and outcome are recorded in the associated SAS server log files. These log files can be processed by the Audit, Performance and Measurement (APM) package, a free tool from the SAS Support website [ref 4]. This package allows a site to monitor the state of the SAS 9 EBI architecture, implement audit authentication reports, and review reports on the performance and usage of the server environment. Figure 4 contains an overview of how the tool collects data. See the Recommended Reading section for papers describing the APM data and its code.

Simplified SAS EBI configuration with Audit analysis

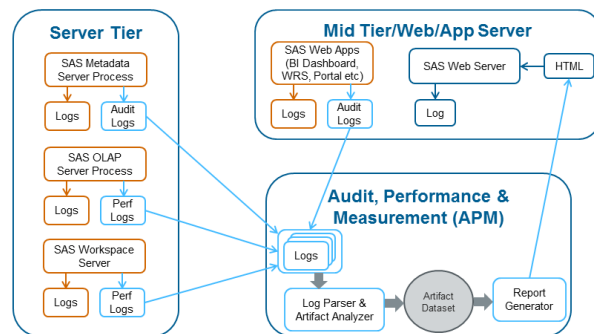


Figure 4 Audit analysis

With the release of SAS 9.4 and the SAS Environment Manager, a web-based administration solution provides SAS administrators with a dashboard that presents the collected data from monitored resources to monitor the health and operation of a SAS environment. For more information regarding the SAS Environment Manager, refer to the *Monitoring Your SAS Intelligence Platform* paper [ref 5].

Understanding the APM Data and Reports

The three main reporting areas that the APM package offers out-of-the-box are:

- SAS Metadata Server audit reporting
- SAS Enterprise BI Server analytic server performance usage reports
- SAS Enterprise BI environment status reports

In Figure 5, you can see examples from the APM home page.

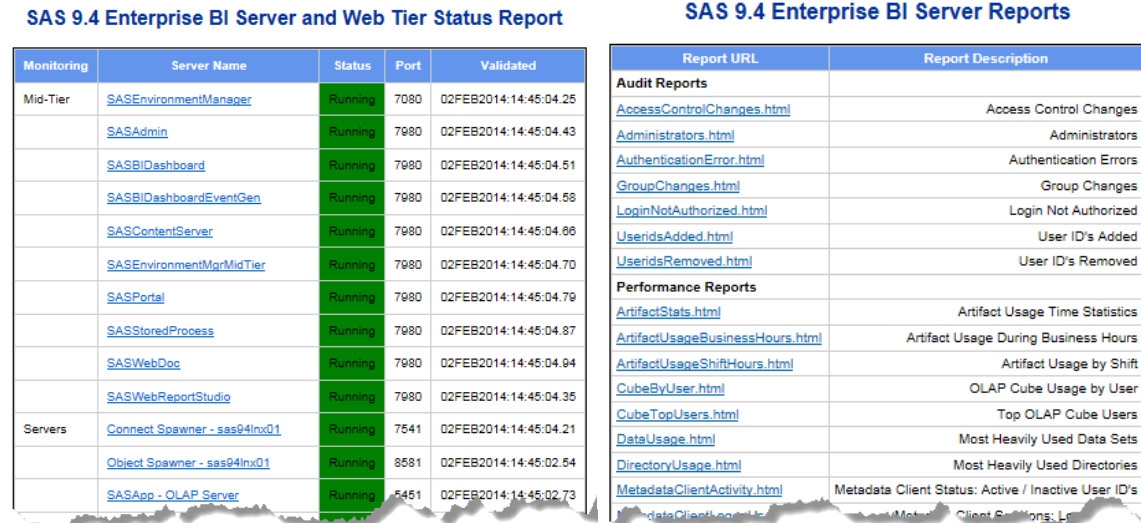


Figure 5 APM Report Examples

Whilst this gives a SAS administrator an insight into the running of their SAS 9.4 EBI platform, it doesn't provide a visual overview of business related metrics with the ability to gain further insight through exploration and advanced analytics that SAS Visual Analytics offers.

Taking Advantage of the APM SAS Code and Data

The APM package installs SAS code on the server to generate tables and reports. And in the process of installing it you schedule it to run regularly so the tables get updated every so often. By taking advantage of this we can use the tables and SAS code in our own SAS Enterprise Guide projects or SAS Data Integration Studio jobs to prepare the data for our dashboard. An "EBIAPM Artifacts" metadata library is created on the server that contains the artifact and audit related tables and a folder structure of \apm94\SASEnvironment\SASCode\Jobs\usageReports contain the SAS code for the APM reports.

Extracting the Data

Once you've examined your business requirements and the metrics needed for your dashboard, you can extract the data from the related APM tables using the ETL tool you are familiar with (SAS Enterprise Guide or SAS Data Integration Studio). The jobs can be scheduled to run after the APM daily log reading and report update has been completed so that the dashboard is updated daily.

Depending on what you require in your dashboard, the data can be sourced from a combination of the APM generated tables, SAS metadata – such as running the MDUEXTR macro (documented in the SAS® 9.4 Intelligence Platform: Security Administration Guide, see the Recommended Reading section). You can collect identity information from the metadata or from a third party tool, such as the Metacoda Plug-ins [ref 6].

Load the Data into SAS Visual Analytics

With the tables converted into SAS data sets, you can load them into SAS Visual Analytics in a variety of ways: Interactive load, autoload, local import, remote import, and data query. For ad-hoc data loads, the local import method might be suitable.

However as we want to ensure our data is loaded after the batch jobs are complete we use the Autoload facility to also accommodate reloading the table too. The Autoload data directory becomes the synchronization point for the source tables. Within the SAS Visual Analytics 6.3: Administration Guide, see the Recommended Reading section,

there are details on how the Autoload facility works, where you need to place your SAS tables and how to schedule the autoload.

CREATING A DASHBOARD

At this point, you should have your goals established, some dashboard layouts, and the associated data loaded into the system. The SAS Visual Analytics Designer is a great way to create dashboards and interactive reports. The tool has indicators, graphic objects, and interactivity that are useful in building a dashboard.

Working with Indicators

Indicators are used to represent metrics. It's not always clear the best way to present the data and you have to try several methods. It's easy with SAS Visual Analytics because you can quickly change the data from one visualization to another.

It helps when you understand the purpose of the metric. For instance, you could determine how many users there are and show how many unique users there was each week. However if you plot this information across a time line you really only understand the login trend. You learn that more users access the system at the beginning of the month and during the holiday shutdown, almost no one used the system. These are all facts, but not conclusions.




Figure 6 Example of a Simple Indicator Layout

You need to understand the purpose of the metric. For example, many executives want access to the system but then rarely use it. Is it likely we are going to remove their login? Maybe the team wants to ensure user adoption. With lack of training, some users revert to storing data in Excel or something similar. The intent of the goal is to ensure 80% of the users are active each month. In this case, you would construct the goal as a percentage (active users this month/total users registered). You could also present the data with an indicator, as shown in Figure 7. Then you could link the indicator to the list of inactive users for follow-up.




Figure 7 Slider Indicator Example

Presenting the Dashboard

As stated the dashboard design is an interactive process. One advantage of the SAS Visual Analytics Designer is that it is very easy to make changes. With the user present, you can create new categories and calculations, place new indicators and rearrange the layout quickly. With the user seeing the live results, it is easy to agree on a final layout and dashboard.

The Visual Analytics Designer also offers more interactive features that may be new to some users. Stephen Harris, who has implemented dashboards for large organizations, noted that user adoption is often an obstacle. He suggests that you work closely with new users during the transition phase [ref 3] by ensuring the user understands how the drop-down filters work and how to click to link to other report sections. He further suggests that you attend meetings where the dashboard is used to observe how the team interacts with it. "This is why your interaction is so vital. If you take the time, you will increase user adoption," Harris said.

A dashboard is an investment. By working together, the two groups are able to monitor activity and resources and become more proactive in identifying issues quickly, recognizing patterns for potential end-user training and

identifying opportunities for resource maintenance, effectively reducing an organizations business risk.

When SAS administrators understand how the tool actually provides value to the organization and the business users understand the background processes required to support the tool, they start to form a common language that makes other projects easier to implement.

SECTION TWO: USING THE ADMINISTRATOR DASHBOARD

The remainder of the paper uses case studies based on the dashboard to see how to integrate your dashboard with the SAS Visual Analytics Explorer and third party tools, such as Metacoda Plug-ins. These case studies show how to explore and determine underlying issues using case studies.

CASE STUDY: MONITORING A SYSTEM ROLLOUT

A business intelligence system usually starts in one area of the company and then spreads to another. As the rollout process occurs, new users are provided training and access. The business sponsor wants to ensure system adoption and be able to address any issues quickly. A dashboard that tracks the rollout activity would help. Our scenario below is to examine the adoption of our 51 users after training completion. Here are the goals associated with our rollout process:

1. Ensure new users have accessed the system and reports within 30 days of training.
2. Monitor the implemented security model to ensure it remains consistent to plan.
3. Monitor resources such as user elapsed CPU time.
4. Monitor user's SAS jobs to ensure they are run during business hours.

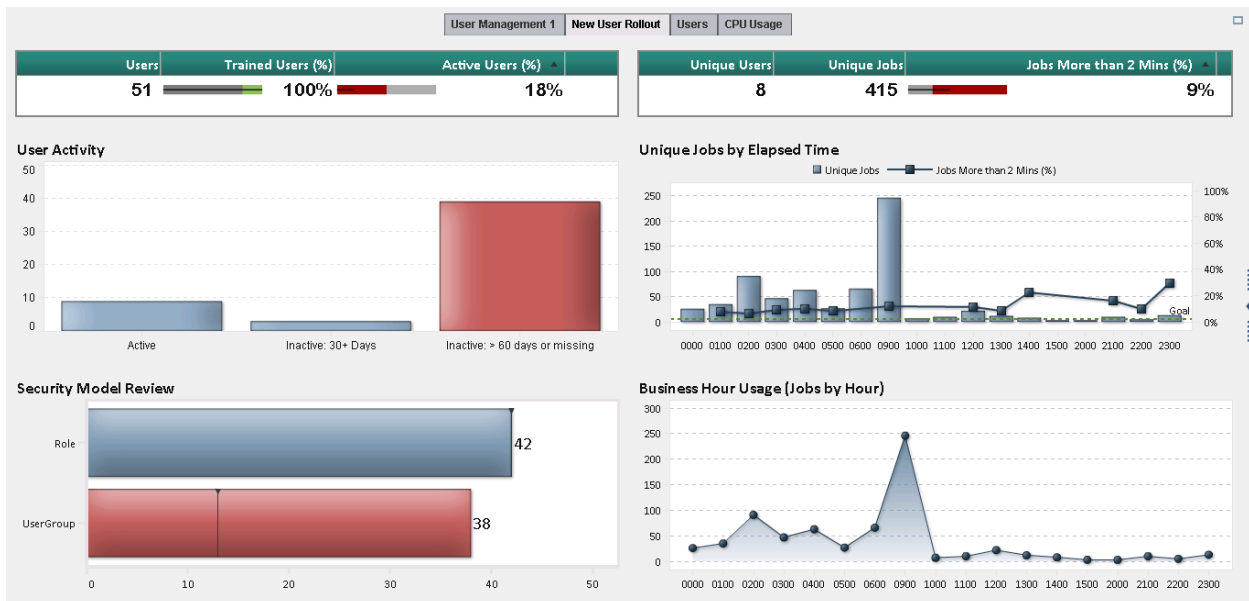


Figure 8 Tool Rollout Dashboard

The dashboard provides a great way to monitor the goals set for optimal business adoption, however it doesn't provide the insight needed to explain and predict usage. As such, we will further analyze the data using SAS Visual Analytics Explorer to investigate the underlying data and identify patterns. Some areas that are of interest based on our goals are:

- Explore server usage measures to identify users that have processes with average CPU times greater than 2 minutes.
- Analyze server usage metrics across the day to understand server usage and user patterns.
- Forecast future user elapsed CPU time to determine whether a goal of less than 2 minutes can be achieved.
- Identify potential SAS training by examining SAS procedure or DATA step usage.

GOAL 1: ENSURE NEW USER TRAINING AND ADOPTION

In the top-left area of the dashboard (shown in Figure 9) there are several monitoring points for user adoption. In this scenario the rollout started 45 days ago. The 51 users were trained in the first week, which is shown with 100% training. However only 15% of the users are active on the system, which makes the indicator red. In the chart immediately below the table, you can see another red bar indicating 39 of the users are inactive or worse, missing, indicating they haven't even logged into the system!

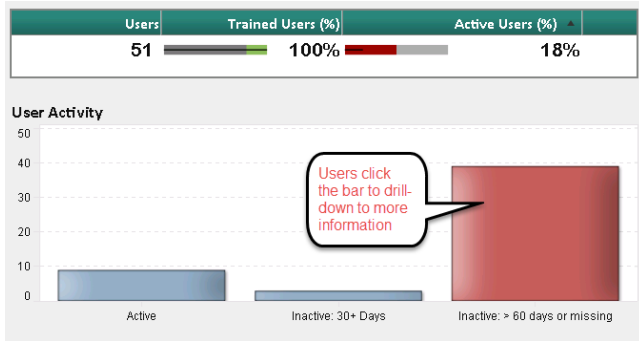


Figure 9 Red indicates underperforming areas

Drill Down for More Information

You can get more information about the underlying data when you click on the red bar. The section linking allows you to go to another report section that contains a summarized user table (see Figure 10) outlining the number of users who are Active and Inactive displays. We have separated Inactive to >60 days or missing and between 30 and 60 days so that we can contact those users that haven't logged in yet (missing) and those that haven't logged in for more than 2 months, collectively.

Incorporating interactivity with a dashboard object is not a standard practice, however the department manager has requested to be able to click on the cell of Inactive to bring up a list of users he can contact. This can be achieved nicely in SAS Visual Analytics through section linking.

name	Active Userid	login_id	Last Login Date
Patrick Paxton			
SAS Administrator		(SASDeployment...	
Lisa Lowes			
Alice Adams		demoalice	.
Koby Kogan (Unrestricted)			.
Bob Baxter		demobob	.
Jane Jacobs (Unrestricted)			.
Eve Evans		demoeve	.
David Doyle		demodavid	.
Fred Faulkner		demofred	.
Carol Charleston		democarol	.
Harry Hanes		demoHarry	.
Ivy Ives		demoivy	.
Jack Jamieson		demojack	.
Kate Knowles		demokate	.
Ian Irons		demoian	.
Hannah Hanes		demohanna	.
Gayle Gordon		demogayle	.
Oliver Oxford		demooliver	.
Natalie Nixon		demonatalie	.

Figure 10 Users who haven't logged in - missing Last Login Date

From the table in Figure 10, we can see from the subset list of users their Last Login Date is missing meaning they haven't logged into the system. The department manager is curious to know why there is no login_id though for Patrick and Lisa compared to the other users. He can see that the special unrestricted accounts for Koby and Jane also have no login_id and is concerned whether this also means Patrick and Lisa have unrestricted accounts? The department manager has a chat with the SAS Administrator who investigates Patrick Paxton and Lisa Lowes accounts using Metacoda Plug-ins within the SAS Management Console.

Use SAS Management Console and Metacoda Software for Extra Information

In Figure 11, the SAS administrator uses Metacoda Plug-ins User Reviewer [ref 6] to investigate. He quickly sees that both Patrick and Lisa don't have a login via the warning indicator in the Login column and also there is no login information for the selected user, Patrick, in the lower pane. Through his knowledge of SAS metadata security he knows that as the SAS identity is disconnected from the operating system identity both Patrick and Lisa will not be able to log in. The SAS administrator recalls these users were manually created which may have been where the mistake occurred.

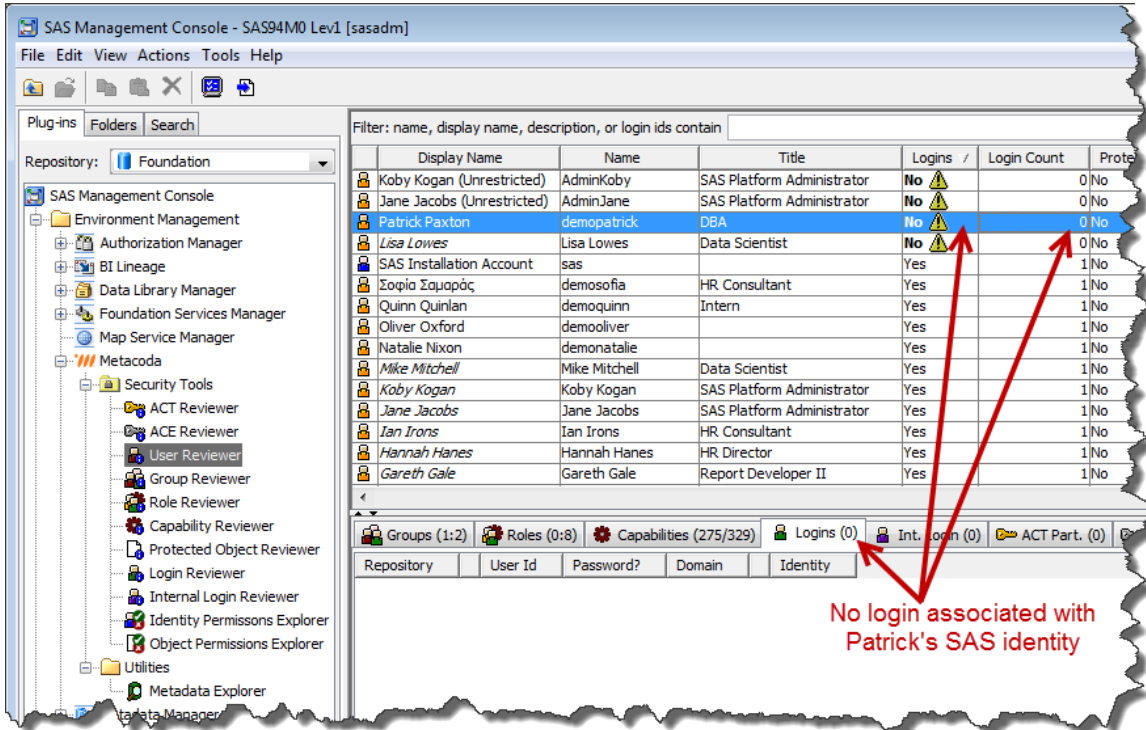


Figure 11 Identifying users with no Login using Metacoda Plug-ins User Reviewer within the SAS Management Console

While selecting Patrick's identity the SAS Administrator quickly verifies other identity related attributes for Patrick. Using the tabs in the lower pane he can see that Patrick is a member of 2 groups implicitly (SASUSERS and PUBLIC) and those groups indirectly provide membership of 8 roles, in turn providing him with 275 capabilities.

GOAL 2: MONITOR THE SECURITY MODEL

One of our other goals is to ensure the implemented security model remains consistent to plan by monitoring the number of user groups and roles. By having a component in our dashboard as in Figure 12, a SAS administrator can quickly see through the display rule if there is a significant change in the number of user groups and roles. This is an indicator to help us identify whether it remains consistent (blue) or no longer consistent (red).

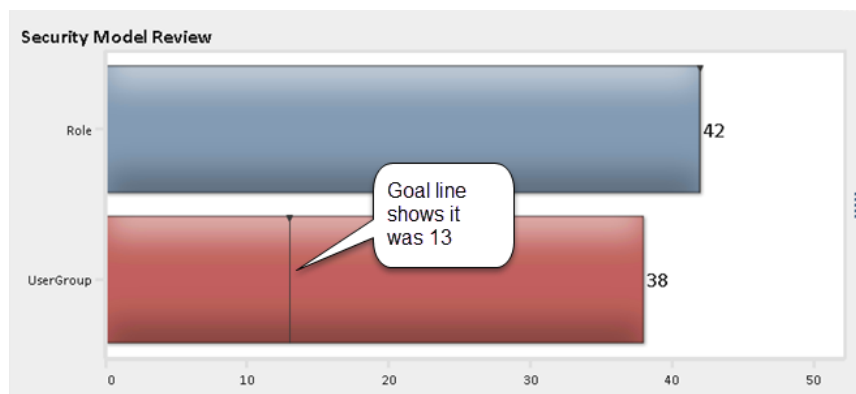


Figure 12 Dashboard component showing user groups and roles

From this component, we can see we have 25 more groups than we should! Our security plan has 13 groups... *what's going on?*

Exploring the Extra Groups

We can click on the chart to view a Groups report as in Figure 13. This report is generated by the Metacoda Security Plug-ins batch export facility [ref 6] and made available on our web server for interactivity, as an external link, with our dashboard component.

Here we can easily see in that there are many groups and from the warning indicator on the Empty Group column we can see these unexpected additional groups with no members. The group icon in the second column has a blue and yellow jacket highlighting that it is an external group rather than an internal group (red and green jacket). This visual indicator helps the SAS Administrator identify that they need to talk to the IT department about Active Directory synchronization.

Row	Repository	Display Name	Protected	ACE Identity	Empty Group	Direct Members	Implicit Members	Logins	Contains Self
1	Foundation	SASUSERS	Yes	Yes	No		No	No	No
2	Foundation	PUBLIC	Yes	Yes	No		No	No	No
3	Foundation	SAS Administrators	Yes	Yes	No	2	No	No	No
4	Foundation	SAS System Services	Yes	Yes	No	1	No	No	No
5	Foundation	SAS General Servers	No	No	No	1	No	Yes	No
6	Foundation	BI Web Services Users	No	No	No	1	No	No	No
7	Foundation	BI Dashboard Administrators	No	No	No	1	No	No	No
8	Foundation	BI Lineage Users	No	Yes	No	1	No	No	No
9	Foundation	BI Dashboard Users	No	No	Yes ⚠	0	No	No	No
10	Foundation	Demo: Contains Self 1	No	No	Yes ⚠	0	No	No	No
11	Foundation	Demo: Contains Self 2	No	No	Yes ⚠	0	No	No	No
12	Foundation	Vegas Enterprises: Oracle Users	No	No	Yes ⚠	0	No	No	No
13	Foundation	Vegas Enterprises: DB2 Users	No	No	Yes ⚠	0	No	No	No
14	Foundation	Vegas Enterprises: Asia Pacific Region	No	No	Yes ⚠	0	No	No	No
15	Foundation	Vegas Enterprises: North American Region	No	No	Yes ⚠	0	No	No	No
16	Foundation	Vegas Enterprises: U.S.A.	No	No	Yes ⚠	0	No	No	No
17	Foundation	User group (external) Canada	No	No	Yes ⚠	0	No	No	No
18	Foundation	Vegas Enterprises: Australia	No	No	Yes ⚠	0	No	No	No
19	Foundation	Vegas Enterprises: New Zealand	No	No	Yes ⚠	0	No	No	No
20	Foundation	Vegas Enterprises: AU/QLD	No	No	Yes ⚠	0	No	No	No
21	Foundation	Vegas Enterprises: AU/NSW	No	No	Yes ⚠	0	No	No	No

Figure 13 Metacoda Security Report: Groups - identifying empty external AD groups

Creating metadata security reports in batch for a SAS environment can also be useful for auditing purposes, as well as providing the ability to compare environments over time.

EXPLORING THE DATA WITH SAS VISUAL ANALYTICS EXPLORER

The SAS Visual Analytics Explorer is a part of the SAS Visual Analytics architecture. It allows you to spot trends in data, find unknown relationships and take advantage of the advanced analytic capabilities from SAS to gain insights and forecasts as described in the blog post, *Look Before You Leap with Data Visualization* [ref 7]. We will now examine our goals further by exploring our data with some further advanced analytics.

GOAL 3: MONITOR USER CPU TIME

On the dashboard, we noticed that the server CPU metrics appeared to be high and not meeting our business objectives. As the new department has new users with introductory SAS skills it was set that processing times should be less than 2 minutes. As the users evolve in their SAS development it would be expected that this objective will change, and by having a dashboard we can monitor this over time too. While we are able to gauge whether processing times are less than 2 minutes in the dashboard it doesn't provide us with the details of that underlying metric – who are the users consuming more resources and for what SAS processes. As the underlying data source for the dashboard is loaded into the SAS LASR Analytics Server, it is also available to explore with SAS Visual Analytics Explorer.

From the underlying APM data collected, we have a number of server usage measurements available to analyze. Using SAS Visual Analytics Explorer, we can start our analysis by dragging and dropping the measurements into an autochart visualization to examine relationships. A correlation matrix is generated automatically as in Figure 14. Here we can see there is a strong relationship between total CPU time and user CPU time, which we would most likely expect to see.

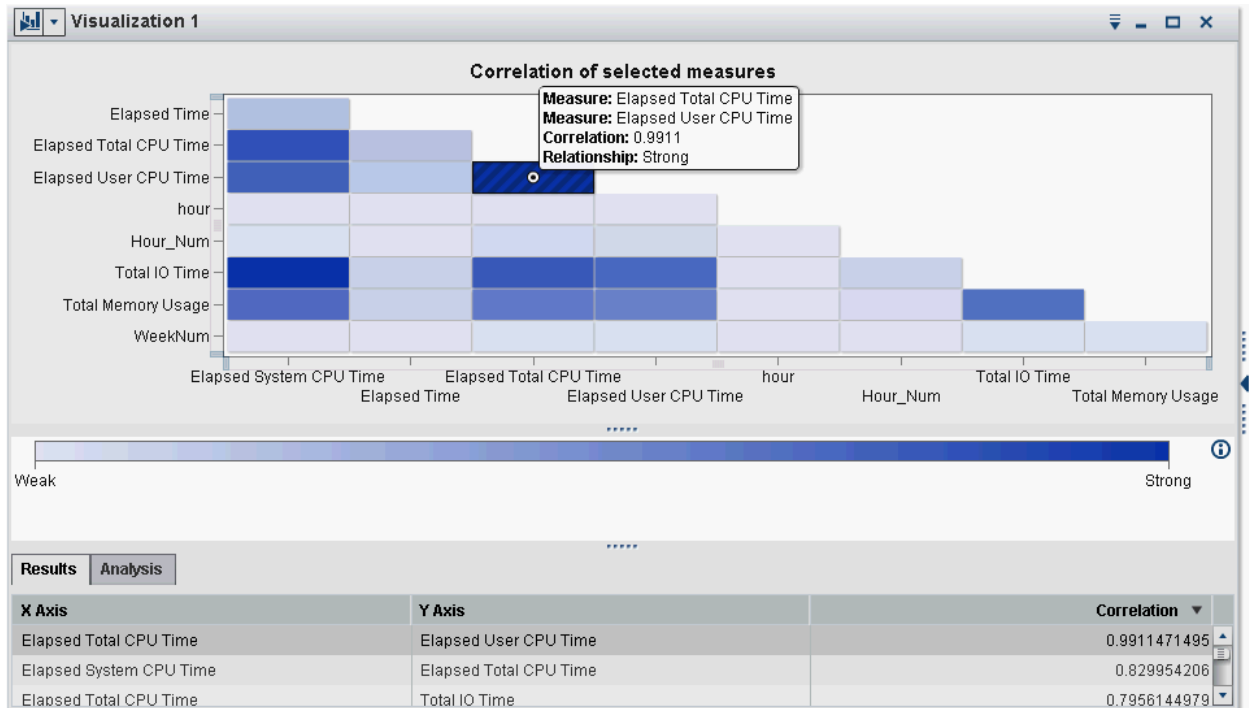


Figure 14 Identify relationships between server usage measurements

Double clicking on the intersection of Elapsed Total CPU Time and Elapsed User CPU Time, we delve into examining the underlying values as a heat map to further examine the relationship. We see in Figure 15 that we have a very strong linear relationship between Elapsed Total CPU Time and Elapsed User CPU Time. What we are immediately drawn to though are two colored cells in the top right corner. As part of our exploration is to determine if there are any users consuming resources for more than 2 minutes (120 seconds) we want to find out who the users are that are consuming CPU resources for almost 20 minutes (1200 seconds) and to know what SAS processes they were running.

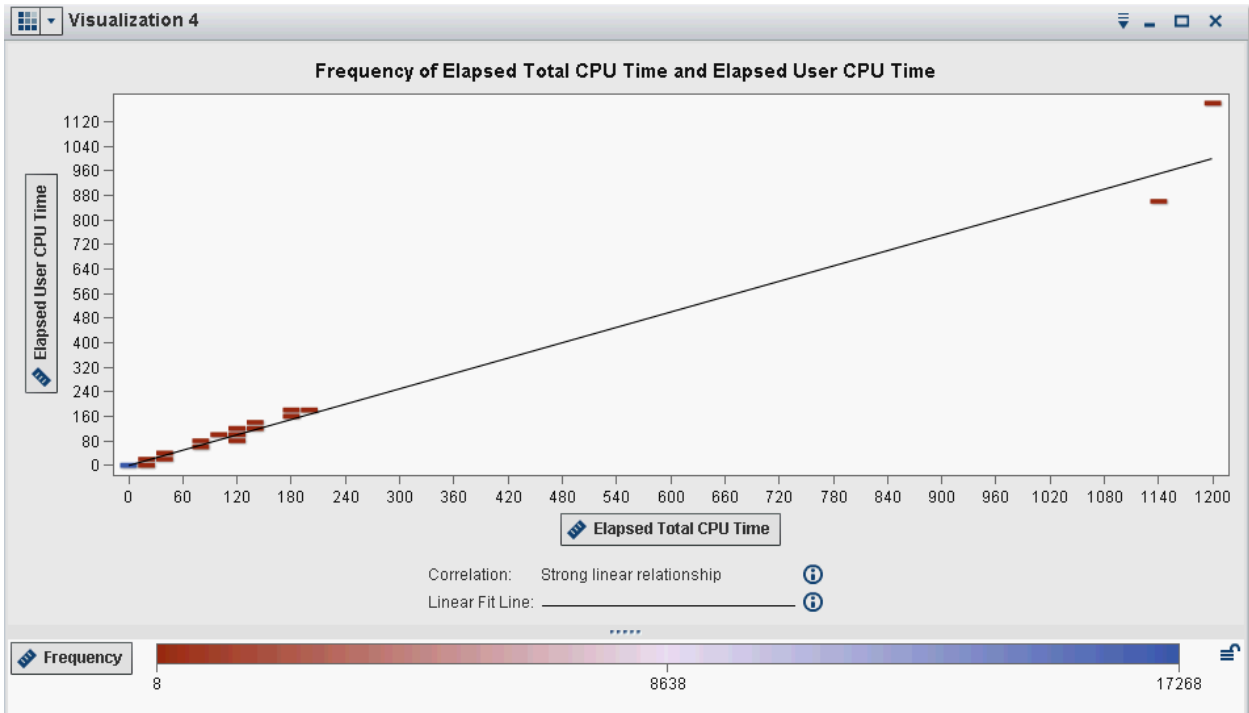


Figure 15 A strong linear trend between Elapsed Total CPU Time and Elapsed User CPU Time

In Figure 15 we also notice in the bottom left corner of the chart the cell color is blue, where all the other cells are red. This indicates we have processes that are low in both Elapsed Total CPU Time and Elapsed User CPU Time however high in the number of processes executed. We would also like to know the user(s) for the blue colored cell.

To identify the users, we alter our visualization replacing Elapsed Total CPU Time with User ID as in Figure 16.

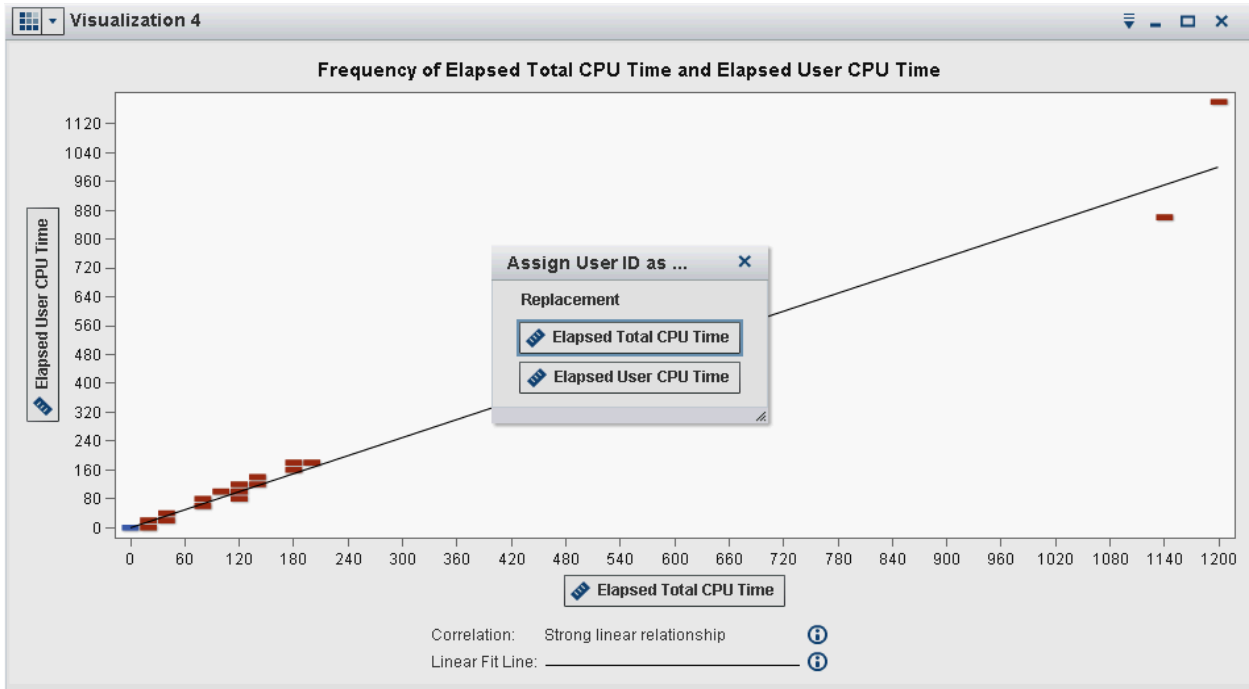


Figure 16 Replacing Elapsed Total CPU Time with User ID

The visualization is refreshed and we can see in Figure 17 that Manager Millie, Mike and Dorothy are the users that are consuming CPU resources more than the other users. And the sas user is the user with the most frequent processes and low Elapsed User CPU Time (blue colored cell).




Figure 17 Identifying users with the highest Elapsed User CPU Time and highest frequency

Show me the stats!

From our analysis so far, we know we have a group of users Manager Millie, Mike and Dorothy who are running processes with Elapsed User CPU time greater than 2 minutes but we don't know how this compares on average with other users. By switching the visualization to a Box Plot in SAS Visual Analytics Explorer we can examine the distribution of Elapsed User CPU Time for Manager Millie. In Figure 18, we see that her average Elapsed User CPU Time (52 seconds) is less than the average 2-minute objective. By further examining this chart, we see that the average for every user (represented by the diamonds in the chart) is below the 2 minute Elapsed User CPU Time which would for these users meet our business goal.

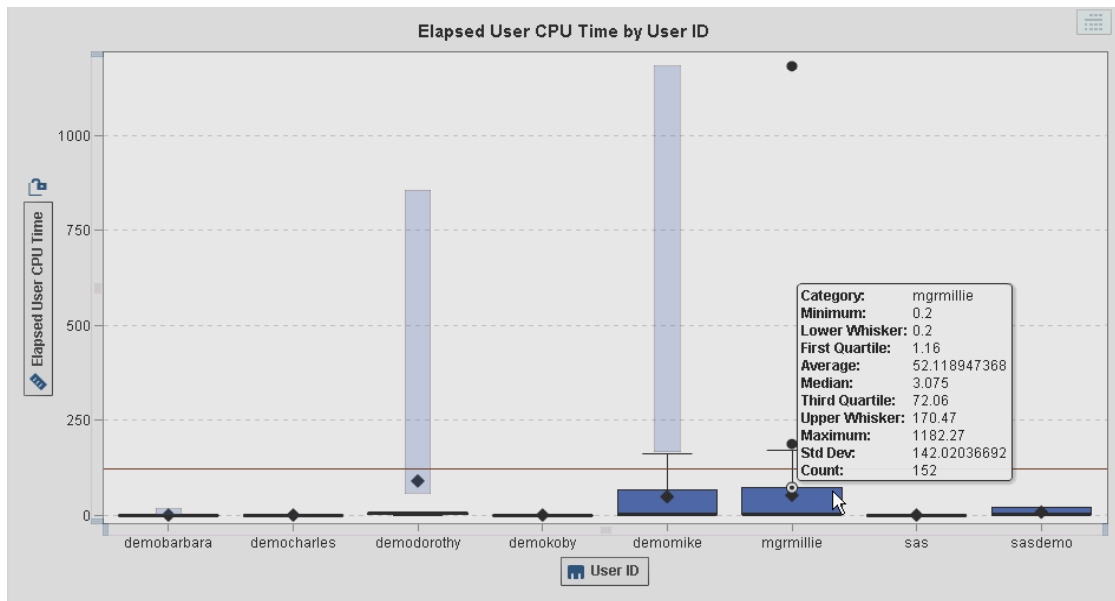


Figure 18 Compare Elapsed User CPU Time distributions

While we have a level of comfort with the average user CPU time being below 2 minutes for our historical data, we are also interested in using this data to help us predict future values, which we will do later with forecasting.

GOAL 4: BUSINESS HOUR USAGE

One of our other business goals is to monitor user SAS jobs to ensure they run during business hours. By changing the visualization to a bar chart and examining the hours of the day, as in Figure 19, we can see Manager Millie runs her processes at the start of the business day. The chart also highlights Mike's elapsed CPU time (yellow bars) and surprisingly we see his processes are taking place early in the morning and definitely not during business hours! We might need to have a chat to him about this as this could affect our nightly backups and other jobs.

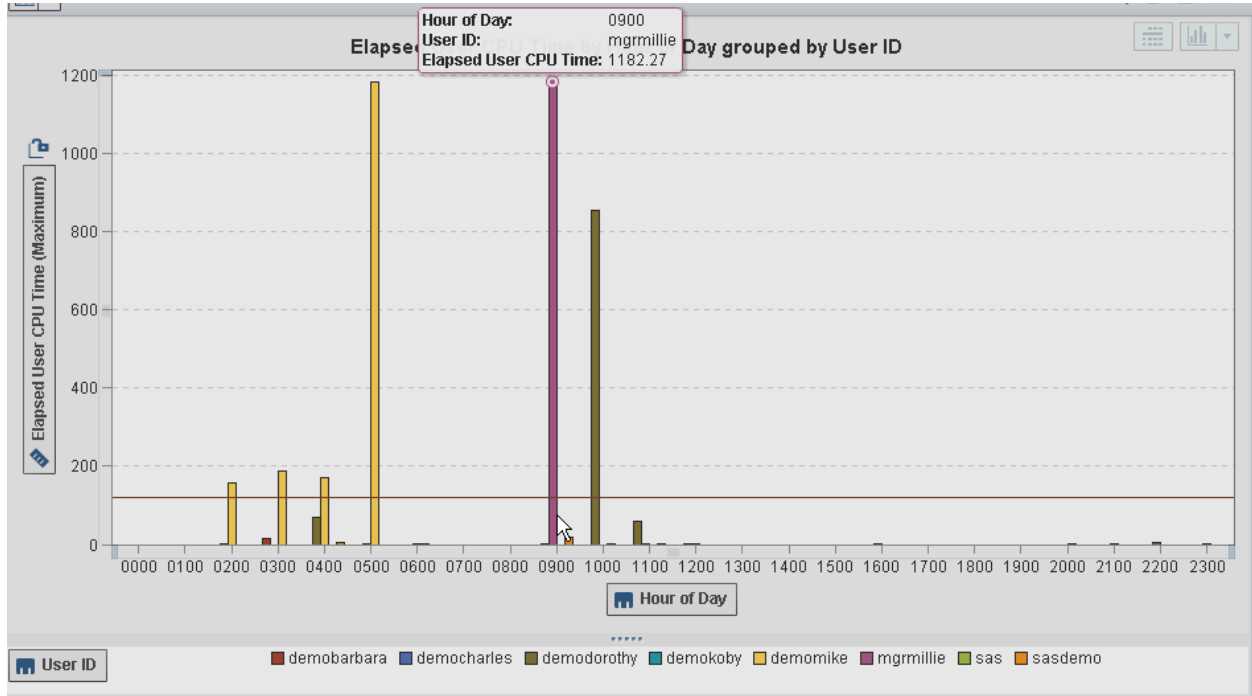


Figure 19 Analyzing elapsed user CPU time by user throughout the hours of the day

In summary, each goal on the dashboard led us down different analysis paths using different tools. When we planned the dashboard (similar to the diagram shown in Figure 3) these explorations would have been how we intended the user to explore the information.

NEXT STEPS FOR YOUR ANALYSIS

Operational and Tactical dashboards, such as the administrative one we created, focus on what did happen. In the following examples, you see the ways that SAS Visual Analytics Explorer can use your existing data to forecast usage.

FORECASTING – CAN WE MEET OUR BUSINESS REQUIREMENT IN THE FUTURE?

In the box plot in Figure 18, we saw the users seem to have an average Elapsed User CPU Time of less than 2 minutes. Perhaps if we forecast based on the maximum Elapsed User CPU Time we can see how well we might meet our business metric of an average of 2 minutes. We can see in Figure 20 there are some spikes, which are probably Manager Millie, Mike and Dorothy's processes based on our initial exploratory analysis above.




Figure 20 Line Chart showing Maximum Elapsed User CPU Time over time

By right-mouse clicking on the chart and selecting forecast, we can easily get a statistical forecast. SAS Visual Analytics adds a line with the predicted values and a colored band that represents the 95% confidence interval as in Figure 21. We also overlaid our line chart with a 2 minute reference line so that we can compare the forecast with our business goal. Based on the wide variability in our historical data the confidence bands are rather large, thus producing a wide confidence interval for our forecast point estimate, 165 seconds. As our business goal is to achieve an average user CPU Time of less than 2 minutes we will do some further analysis to find out what processes they are running and to help them tune the underlying SAS code sent to the server.

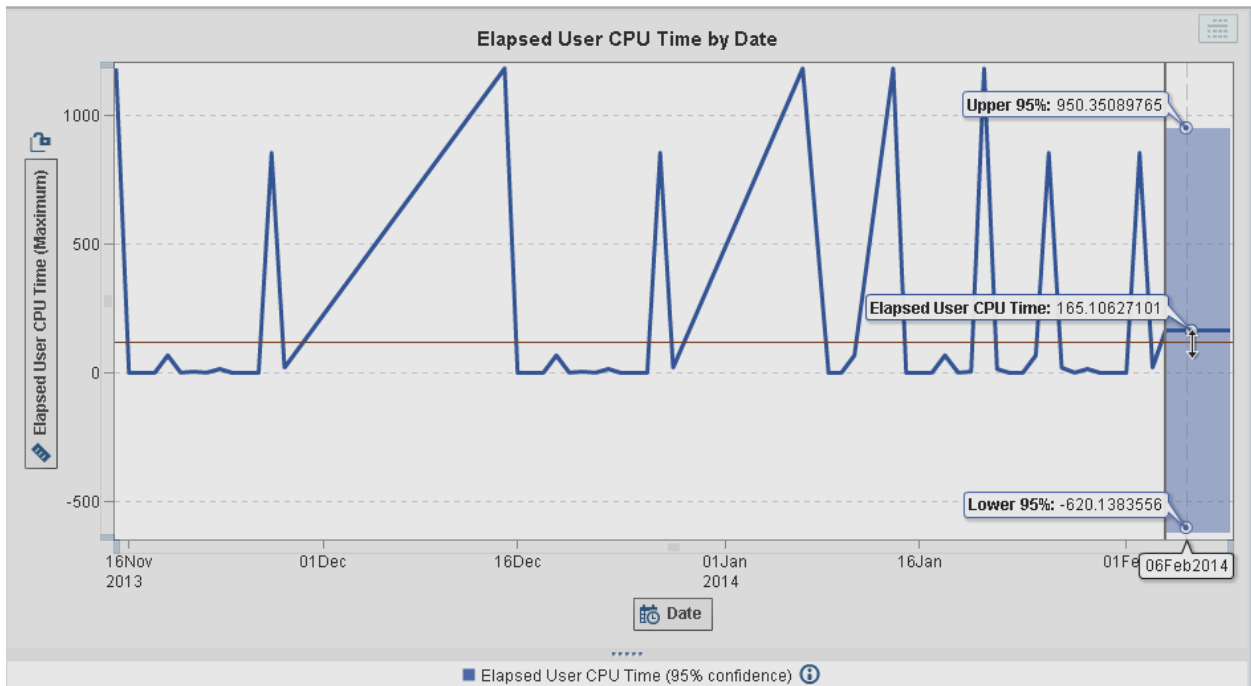


Figure 21 Producing a forecast based on the statistical trends in the data

MAKING PREDICTIONS WITH DECISION TREES

In Figure 19 we analyzed the Elapsed User CPU Time by user by the hours of the day. In Figure 21 we examined the maximum Elapsed User CPU Time by time to get a forecast based on our historical data. With SAS Visual Analytics, we can use a Decision Tree visualization to use the values of our data items to predict the values of Elapsed Max CPU Time.

A decision tree displays a series of nodes as a tree, where the top node is the target data item, Elapsed Max CPU Time and each branch of the tree represents a split in the values of a predictor data item. We will use Elapsed User CPU Time, Hour of Day, Max Memory Usage and Elapsed Time.

In Figure 22, we see that the best data item to explain Elapsed Max CPU Time is Elapsed User CPU Time with a split value of 591.21 seconds. Interestingly for processes where the maximum elapsed CPU time is more than 591.21 seconds no other data item best explains Elapsed Max CPU Time. This seems to suggest that if the Elapsed User CPU Time was more than 5 minutes (600 seconds) then it could be used as a predictor for Elapsed Max CPU Time. Whereas, if the elapsed user CPU Time is less than 591.21 seconds then the hour of the day also best explains it with 0300, 0400, 0900 having similar characteristics to explain Elapsed Max CPU Time.

As you can see, SAS Visual Analytics Explorer provides us an interactive way to generate decision trees to depict graphically likely outcomes on our target data item, Elapsed Max CPU Time.

Over time, these predictors may not be the best to predict Elapsed Max CPU time and as business users evolve in their use of the SAS EBI environment, other goals and therefore target data items are become the focus. As such, using an analytics tool such as SAS Visual Analytics Explorer allows a user to monitor changes to the model through quick data discovery and exploration.

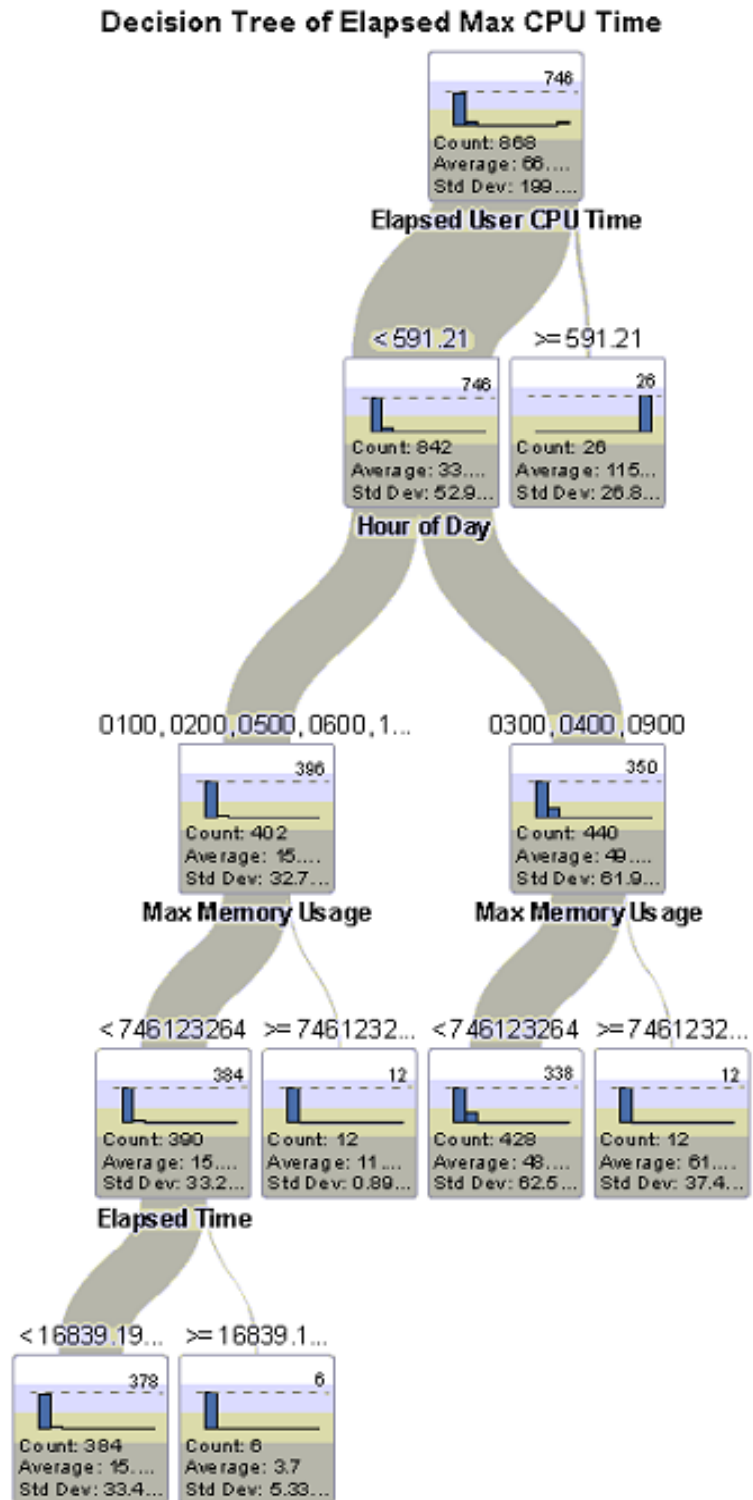


Figure 22 Elapsed User CPU Time is the best data item to explain Elapsed Max CPU Time

OTHER WAYS TO ANALYZE DATA – TAG CLOUDS

In our earlier analysis on Elapsed User CPU Time, we were curious to know what processes Manager Millie, Mike and Dorothy were running to be taking up almost 20 minutes of CPU time – perhaps they need some training? In our dashboard, we have an object that summarizes Elapsed User CPU Time based on the SAS procedure or DATA step process that runs the task. We are able to provide this dashboard summary as it is part of the data collection from the APM logs, see Figure 4. Each request to the SAS server, across all SAS products, whether via SAS code, point and click in SAS Enterprise Guide or SAS Enterprise Miner, the names of the underlying procedures or DATA steps submitted, are logged and collected in APM artifact tables.

We could produce a summary table showing the most used procedure names including DATA step however, that would not be very appealing, and we only get a sense of magnitude when we sort in descending order. An alternative approach would be to investigate the words of the procedures or DATA step in a Word Cloud visualization. In doing so, the size of each word in the cloud will indicate the frequency of the word and we can also use a measure to see the value of that measure with color. As the text field in the collected data is a 1 word procedure or data step we have used the Word Cloud visualization with a category item, rather than a document collection role, assigned to visually identify which procedure name or data step consumed the highest Elapsed User CPU Time and Total Memory Usage, see Figure 23.

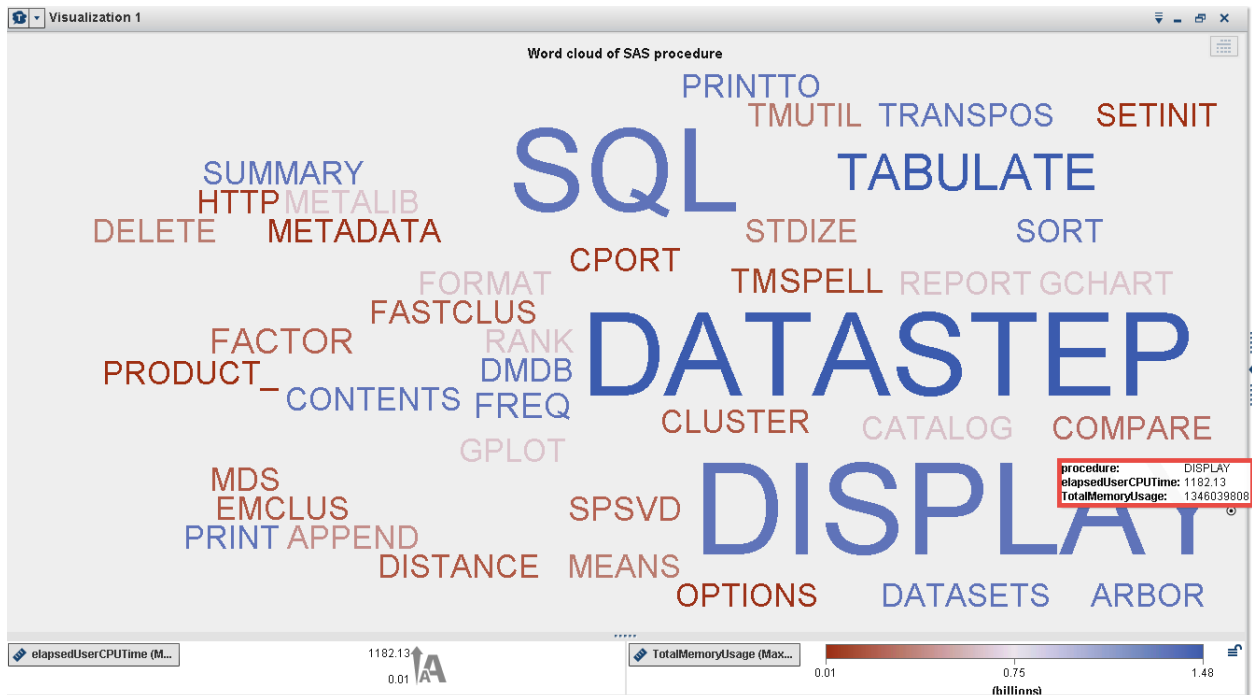


Figure 23 Visualizing the procedure names (and datastep) used against Elapsed User CPU Time and Total Memory Usage

Analyzing the executed procedure names and DATA step with Elapsed User CPU Time (text size) and Maximum Total Memory Usage (text color), we can see that most of the memory and CPU usage being used within DATA step processes, PROC DISPLAY (that SAS Enterprise Miner executes behind the scenes) and PROC SQL. We see this quickly as they are the largest text in blue. Whereas PROC SETINIT, PROC OPTIONS, and PROC METADATA are the smallest text in red meaning the least Elapsed User CPU Time and Total Memory Usage.

As business users evolve with their new SAS environment, it is likely that change requests and issues will be raised for a future phase in the project. It's also possible to integrate data from a service request system with our dashboard to provide metrics on the adoption. As well as providing these metrics, additional insights may also be gained by analyzing the service request text field data in a Word Cloud visualization to identify topics and more commonly reported issues/areas of concern. This quickly identifies issues raised by end users by performing text analytics on the service request free form fields, which may remain undiscovered for analysis only done on service request pre-defined category fields.

From our previous explorations, we know that Manager Millie, Mike and Dorothy are the users with high Elapsed User CPU Time. By quickly creating a bar chart on procedure, grouped by user for the top 10 ranks of maximum Elapsed User CPU Time we can see in Figure 24 that it is indeed Manager Millie & Mike consuming resources with running DATA step and PROC DISPLAY code and Manager Millie with PROC SQL. As the new users haven't yet received their training in SAS Enterprise Miner or SAS Enterprise Guide Query Builder (which uses PROC SQL), we may need to work with the business to ensure that happens sooner if they want to keep to their set business objectives.

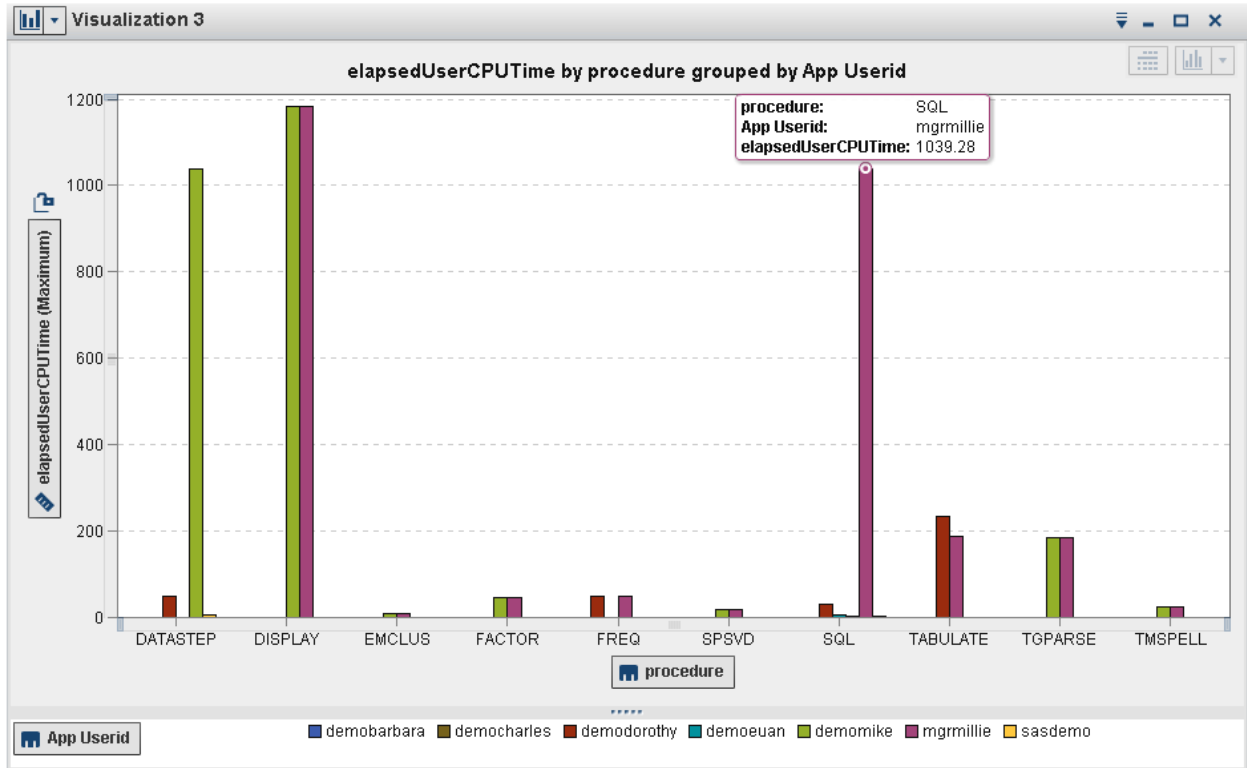


Figure 24 Examining maximum Elapsed User CPU Time by procedure and user

CONCLUSION

By building a SAS Visual Analytics dashboard, the business and SAS administrators can work co-operatively to monitor and manage the user activity, SAS metadata and system performance. We have shown how this can be achieved using the APM generated data that is created from logging within your SAS EBI platform. Furthermore, insights and forecast can be gain by using SAS Visual Analytics Explorer, which can help management become proactive, rather than reactive, to resource related issues.

The dashboard process shown in this paper was applied to SAS administration, but you could use these same steps to create any kind of dashboard for your organization. The key concepts were:

1. Determine what the user wants to measure.
2. Create sample layouts and build multiple prototypes toward the final design.
3. Collect the supporting data.
4. Create the final dashboard.
5. Observe the users interact with the dashboard to ensure adoption.

REFERENCES

- 1 La Grouw, Gail. 2012. *Effective Dashboard Design: Design Secrets To Getting More Value From Performance Dashboards*. Scotts Valley, CA : CreateSpace Independent Publishing Platform.
- 2 Eckerson, Wayne. Profiling Operational Dashboards, Inside Analysis. January, 2012. Available at <http://insideanalysis.com/2012/01/profiling-operational-dashboards/>.
- 3 Aanderud, Tricia. E-Chat Tomorrow: How to Get Users to Embrace Dashboards -- & Other Tips. All Analytics. March 2012. Available at: http://www.allanalytics.com/author.asp?section_id=1858&doc_id=241391.
- 4 SAS® 9.4 Audit, Performance and Measurement Tool. Available at: <http://support.sas.com/rnd/emi/APM94/index.html>.
- 5 Peters, A, Bonham, B, and Li, Z. 2013. Paper 463-2013, *Monitoring 101: New Features in SAS 9.4 for Monitoring Your SAS Intelligence Platform*. Cary, NC: SAS Institute. Available at: <https://support.sas.com/resources/papers/proceedings13/463-2013.pdf>.
- 6 Metacoda Plug-ins <http://www.metacoda.com/>
- 7 Homes, Michelle. *Look Before You Leap with Data Visualization*. March 2013. Left of The Date Line SAS Blog. Available at: <http://blogs.sas.com/content/anz/2013/03/06/look-before-you-leap-with-data-visualisation/>

RECOMMENDED READING

APM DOCUMENTATION AND PAPERS

- *Audit, Performance and Measurement Package for SAS® 9.4 Enterprise Business Intelligence*, Last revised August 16, 2013, Available at: <http://support.sas.com/demosdownloads/downloadfile/114383/APM94.pdf>
- Nelson, Gregory. *Understanding SAS® auditing and performance management*. SAS User Groups Blog. December, 2012. Available at: <http://blogs.sas.com/content/sgf/2012/12/26/understanding-sas-auditing-and-performance-management>.
- *SAS® 9.4 Intelligence Platform, Security Administration Guide*, Cary, NC: SAS Institute. Available at: <http://support.sas.com/documentation/cdl/en/bisecag/67045/HTML/default/viewer.htm>
- *SAS® 9.4 Intelligence Platform: Security Administration Guide, Second Edition*, Cary, NC: SAS Institute. Available at: <http://support.sas.com/documentation/cdl/en/bisecag/67045/PDF/default/bisecag.pdf>
- *Extending the Metadata Security Audit Reporting Capabilities of the Audit and Performance Measurement Package*. October 2010, Cary, NC: SAS Institute. Available at: <http://support.sas.com/rnd/emi/docs/Ebiapm92.Metadata.Security.Audit.Reporting.pdf>.
- *SAS® 9.2 Enhanced Logging Facilities*, 2008. , Cary, NC: SAS Institute. Available at: <http://support.sas.com/rnd/papers/sgf2008/logging.pdf>
- Crevar, M and Brown, T. 2012. Paper 366-2012 *Logging 101: Leveraging the SAS® Enhanced Logging Facility*, Cary, NC: SAS Institute. Available at: <http://support.sas.com/resources/papers/proceedings12/366-2012.pdf>.

VISUAL ANALYTICS

- *SAS® Visual Analytics 6.3: Users Guide*, Cary, NC: SAS Institute. Available at: <http://support.sas.com/documentation/onlinedoc/va/index.html>.
- *SAS® Visual Analytics 6.3: Administration Guide*, Cary, NC: SAS Institute. Available at: <http://support.sas.com/documentation/onlinedoc/va/index.html>.

DASHBOARDS AND VISUAL DISPLAY OF DATA

- Few, S. 2013. *Information Dashboard Design: Displaying Data for At-a-Glance Monitoring, Second Edition*. Burlingame, CA, Analytics Press.
- Tufte, E. 2001. *The Visual Display of Quantitative Information*. Chesire, CT: Graphics Press.

BLOGS RELATED TO TOPIC

- BI Notes for SAS Users, Available at: <http://www.bi-notes.com>

- PlatformAdmin.com, Available at: <http://platformadmin.com/>

ACKNOWLEDGEMENTS

The authors would like to thank their spouses Ken Aanderud and Paul Homes for their continuous support with answering our questions at any time of the day; to Dan Lucas and Brandon Kirk for their collaboration and advice; to Donna Bennett for her suggestions and advice and Rene Cherny for her editing prowess. Any errors or omissions may be credited to the authors.

CONTACT INFORMATION

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

Tricia Aanderud
And Data Inc
tricia.aanderud@and-data.com

Michelle Homes
Metacoda
michelle.homes@metacoda.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.