

Paper 277-2010

## Zen and the Art of Enterprise SAS<sup>®</sup> BI Server Deployment

Dave Steven, Epistemata, LLC, Seattle, WA

*Those who know don't tell and those who tell don't know.*

*-Zen Proverb*

### ABSTRACT

Sales presentations invariably demonstrate vendor CRM software products effortlessly solving all problems with a simple click of the mouse. So why does it always seem to be a nightmare of delays and cost overruns to deploy these products in the real world? The SAS<sup>®</sup> BI (or SAS<sup>®</sup> EBI) Server is no different. Indeed, the superior level of flexibility and customization built into SAS<sup>®</sup> software may seem to complicate matters. The IT Department may incorrectly—and disastrously—apply database engineering concepts when architecting SAS<sup>®</sup> servers, and your local SAS<sup>®</sup> programming Guru really has no background in enterprise systems architecture to be of much help. The way out of SAS<sup>®</sup> BI Server deployment suffering includes a correct understanding of the SAS<sup>®</sup> Intelligence Platform based on experience, a technical appreciation of how it fits into the existing IT infrastructure, familiarity with end-user challenges, and a good measure of tenacity. Specific technical details and managerial recommendations are discussed.

### INTRODUCTION

*When you seek it, you cannot find it.*

*-Zen Proverb*

I can see by the clock in the lower right corner of my monitor, without taking my hand from the mouse, that it is eight-thirty in the morning. The server, even under a load of just three user connections, is thrashing. When it's straining at eight-thirty with just three users, I'm wondering what it's going to be like in the afternoon...

Nothing causes managerial suffering like falling face-flat on the first step of a thousand miles' journey into a new business intelligence deployment. Yet this happens with more frequency than one would expect. Why? Deploying SAS<sup>®</sup> BI (or SAS<sup>®</sup> EBI) Server is quite unlike deploying almost any other sort of application. It's not "plug-n-play" the way office productivity software is, it functions in ways that are technically quite different from database servers like Oracle, Teradata, or SQL Server, and it comes with data management capabilities and enhanced connectivity features unseen in other analytic software packages. Your staff of SAS<sup>®</sup> programming experts is neither trained nor experienced in enterprise-level IT architecture and engineering, and the IT Department is unfamiliar with the technical quirks and caveats that are unique to a SAS<sup>®</sup> BI installation—to say nothing of the peculiar demands of an *ad hoc* analytic environment. And neither typically has experience in managing analytic data and users within the SAS<sup>®</sup> Metadata Server connected component system. Unfortunately, merely corralling personnel from each of these departments into a SAS<sup>®</sup> BI Server planning committee no more produces the necessary expertise than does combining wood with soil produce leaves.

*The obstacle is the path.*

*-Zen Proverb*

Every SAS<sup>®</sup> BI Server deployment is a custom deployment. Let me say that again: *Every SAS BI Server deployment is a custom deployment.* There is no easy formula for getting it right. Any SAS<sup>®</sup> BI Server deployment checklist will include multiple entries entitled, "Figure out how best to deal with ---". The IT infrastructure and policies of your department may or may not be friendly to what needs to be done to ensure a successful SAS<sup>®</sup> BI Server deployment. In some cases, there's a way make an accommodation for the existing infrastructure and to compromise to work within the existing policy. At other times, skillful negation techniques may need to be employed. Perhaps it will be necessary to find a way live without a certain feature available in the software because it's just not supportable by your organization at this time. In any case, remember this: *Every SAS<sup>®</sup> BI Server deployment is a custom deployment.*

The SAS<sup>®</sup> BI Server is what it is, whether you realize it or not. But if it is the right tool for you, then it's a tool worth understanding. "When walking, walk. When eating, eat." Not only does this mean to focus on the task at hand, but to do the right thing for the right activity. The SAS<sup>®</sup> BI Server comprises enterprise-level aspects of statistical analytics, graphing and visualization, relational database server, data warehouse management, data parallel processing, enterprise system management, grid computing, web server, OLAP, ETL—and a whole bunch of various and sundry other acronyms. A holistic system requires a holistic approach when designing and building it.

One would no sooner eat with ones feet as to walk on one's teeth. Yet it's not terribly unusual to see folks attempting to architect and deploy SAS<sup>®</sup> BI Server using an inappropriate approach or resources not familiar with the obstacles, hurdles, and demands of deploying and managing a SAS<sup>®</sup> BI Server in an ongoing production environment. If you've decided that SAS<sup>®</sup> BI is the right resource for your staff, you'll also need to identify the right resources to deploy and maintain your SAS<sup>®</sup> BI Server environment. And the earlier, the better.

## PLANNING

*Should you desire great tranquility, prepare to sweat white beads.  
-Hakuin*

In the immortal words of SAS<sup>®</sup> BI Server Zen Masters Mike Vanderlinden and Brian Varney, "*The most important aspect of a [SAS<sup>®</sup>] BI implementation is planning.*" Truer words have never been spoken. The upside to having informed and experienced planning as prelude to deploying the SAS<sup>®</sup> BI Server is truly a thing of beauty. The penalty, by contrast, for failing to execute effective SAS<sup>®</sup> BI Server planning is... well, let's just say it's not a Zen Experience. Spring will *not* come, and the grass will *not* grow if you sit peacefully doing nothing during the initial planning stages. To get a Zen Experience out of your SAS<sup>®</sup> BI Server, you'll need to put some "Zen" into the planning of its design, installation, deployment, and maintenance. Robert Pirsig, author of *Zen and the Art of Motorcycle Maintenance*, observed that "The only Zen you find on the tops of mountains is the Zen you bring up there." The climb is steep and strenuous, but if you bring enough Zen with you to your SAS<sup>®</sup> BI Server deployment planning, the view can be spectacular. Out-of-the-box, SAS<sup>®</sup> BI Server functions well. Everything that comes beyond that is left to the SAS<sup>®</sup> BI Architect/Administrator to configure according to the needs of the department(s) for which it was deployed and *according to level of his or her own expertise*. SAS<sup>®</sup> Institute wisely designed its products to be extremely customizable. The best SAS<sup>®</sup> BI Server installations are the ones with the best—and best documented—customizations such as job control, user-level disk usage monitoring, automated SASWORK "scratch" space maintenance, and custom client error logging to enhance technical support for SAS<sup>®</sup> analysts and SAS<sup>®</sup> Enterprise Guide<sup>®</sup> programmers,

Not only does the critical planning stage involve scoping-out physical resources such as available RAM (typically *far* more important for SAS<sup>®</sup> servers than CPU speed), number of CPU nodes (should be at least one more than the usual number of concurrent non-threaded SAS<sup>®</sup> jobs), network bandwidth, RDBMS connectivity, disk space, I/O channels, and the like, but also intangible assets such as security requirements—including individual and group data access levels, collaboration practices, and file sharing requirements; the desired level of access to individual SAS<sup>®</sup> logs for technical support considerations (See for example the discussion regarding the use of Unix ACLs in the section on "The Dharma of the SAS<sup>®</sup> BI Server" below), and whether SAS<sup>®</sup> BI Server software patches and upgrades will be handled in serial or parallel—or some hybrid of the two.

Imagine, for example, this actual scenario I once came across: After gathering business requirements from the analytic staff, working closely with the IT Department and going through the pre-installation checklist, the SAS<sup>®</sup> BI Server software is installed and initially configured—only to find that the security policy for accessing the server from desktop clients prohibit the SAS<sup>®</sup> Integrated Object Model (IOM) communications protocol. This isn't something that would be familiar to even the most adept SAS<sup>®</sup> statistical modeling guru, and the likelihood that anyone in the IT Department has even heard of the SAS<sup>®</sup> IOM is floating close to zero. The installation engineer from SAS Institute, not being familiar with what's "under the hood" with respect to specific security protocols at your company confirmed that an administrative login was available and that the correct server ports and services were officially cleared. Obtain an exception to the security policy and waiting to have it implemented at this point can set back the deployment schedule by weeks or, more likely, months.

Even with system security issues in the clear, it's well worth the time and effort to plan early for how *data* security and access requirements may evolve over time, as the SAS<sup>®</sup> metadata Access Control Templates (ACTs) are hierarchical and complex and offer configurable—hence convoluted—concepts such as least permissions inheritance. While minor access modifications are relatively trivial to calculate and deploy, major overhauls in access definitions typically

involve invoking a software development lifecycle (SDLC) process including development and test environment procedures along with the requisite process change notifications, user acceptance testing (UAT) and an announced “changeover” schedule. Effective initial planning can make the trivial amendment process more likely while poor initial planning can easily result in the need to invoke an SDLC process much more frequently for what should have been a trivial modification.

## PERSPECTIVE

*Move and the way will open.  
-Zen Proverb*

A common misstep is to approach deploying a new SAS<sup>®</sup> BI Server as one would deploy a new database server. This would be a good approach except for the facts that SAS<sup>®</sup> tables generally aren’t organized in the same way that typical database tables are organized; analytic data generally are accessed and manipulated in ways not usually seen among database applications; and the statisticians and programmers developing solutions using SAS<sup>®</sup> BI Server generally do not operate in ways typically observed among the familiar users of database servers. Even when data are accessed by SAS<sup>®</sup> from a RDBMS using, say, SAS<sup>®</sup> Enterprise Guide<sup>®</sup>, consideration must be taken for managing the number of open RDBMS connections and how the SAS<sup>®</sup> System—and SAS<sup>®</sup> users—make use of scratch space (i.e. “SASWORK”). There is more than one way to do this, depending on, among other things, how your department is organized.

Don’t merely count on your IT department, SAS programmers, and the implementation engineer working with your SAS Institute Account Representative. Generally IT departments are unfamiliar with how SAS<sup>®</sup> behaves and is used; SAS programmers are unfamiliar with enterprise-level IT requirements such as the SDLC, change control, documentation, and multi-user environment issues; and the implementation engineer from SAS Institute typically does not have sufficient access to the detailed business processes and relationships sufficient to define very much more than the focus of his or her role: a suggested functional installation plan and proof of concept. SAS<sup>®</sup> BI software, on the other hand, goes well beyond “functioning;” it’s all about performance, efficiency, and generating nearly any sort of results needed. And this is why it behooves you to ensure that you’re getting what you paid to have from your SAS<sup>®</sup> BI Server deployment.

It’s not unusual, however, to receive assurances from various experts, adamantly insisting that they can deliver what you want. This may come from within your department or enterprise, or it may come from a consultant or representative you’ve hired. And it typically comes from folks who have proven themselves to be talented—even attaining “guru” status—for what appears to be closely related functions, such as SAS programming. How does one determine, through the cacophony of technical jargon, acronyms, and unintelligible geek-speak, who is truly knowledgeable about SAS<sup>®</sup> BI Server design and administration vs. who merely *thinks* they’re knowledgeable?

It’s really not as difficult as it may seem: The person you’re looking for is the one who can (1) pass an interview with your IT Department for a position as a Junior Systems Engineer, and (2) pass an interview with your SAS programming staff for a position as a Senior SAS Programmer. Interview them for both. If you’re paying attention to just one or the other, you’re likely to be in for some displeasure somewhere down the line. It also helps if this person can produce a certificate documenting completion of SAS Institute’s course on Platform Administration. And it would be a huge plus for this person to have had a few years’ experience administering SAS<sup>®</sup> BI Server for a department of analysts. It’s one thing to get SAS<sup>®</sup> BI Server online and functioning; it’s a whole other thing to customize it, maintain it, and modify it do what your staff needs it to do.

Every SAS<sup>®</sup> BI deployment is a custom deployment.

## THE DHARMA OF THE SAS<sup>®</sup> BI SERVER

*From the withered tree, a flower blooms..  
-Zen Proverb*

What is the solution to avoid becoming *Lost* in the attempt to make sense out of deploying SAS<sup>®</sup> BI Server by embracing the traditional approach—an approach which combines IT staff with business/analytic staff in a dualistic

cacophony of personnel who really only possess fragments that hint at the answers everyone is striving to understand? Zen Masters of the SAS® BI Server typically employ the Zen concept of “non-duality” to intertwine the enterprise IT requirements along with the business/analytic requirements into a cloth of its very own unique pattern. Of the many threads which make-up that cloth, here we will briefly address SAS® BI Server disk space management and SAS® Enterprise Guide® programming technical support.

IT departments typically monitor disk usage at the server or SAN allocation level, responding to alarms when aggregate usage passes a certain threshold. This approach, while effective for most database server applications, very typically is either too sensitive or too insensitive for most applications employing the SAS® BI Server. Aggregate disk usage alarms set at too high a level usually trigger too late to avoid disk usage bottlenecks that delay business results by hours or days. But there is little room in this aggregate approach before the level is too low, resulting in disk usage alarms that trigger every few minutes.

Zen Masters of the SAS® BI Server pay less attention to aggregate disk usage levels in preference to *user-level* disk monitoring, employing scripts that, for example, run a Unix “du” command script as root. Temporary and scratch space usage, both within a “sandbox” as well as on production systems, for the majority of SAS® and SAS® BI Server applications, typically generate usage spikes that may endure for seconds, minutes, hours, or even days. Disk usage management on a SAS® BI Server, to be effective, requires an in-depth understanding of the “culture” of how the machine is used, and what the “usage profile” of individual analysts looks like.

One of the most popular applications of SAS® BI Server is to support analysts and statisticians using SAS® Enterprise Guide®, which leverages centralized data sharing and management, a visual programming development environment, stored process support, and OLAP cube analytics, among other things. One of the apparent caveats of deploying SAS® Enterprise Guide® is the apparent ephemeral nature of the SAS® logs that Enterprise Guide® generates. It's only too common that an enterprising SAS® Enterprise Guide® user runs into a problem for which the natural human response seems to be to exit Enterprise Guide® and restart... and then call for technical support. But what happened to the error log that never seemed to be generated in Enterprise Guide® and/or is no longer there once they restarted their Enterprise Guide® client?

“Out of the box,” the SAS® BI Server is not pre-configured to generate Enterprise Guide® logs on the server persistently. But it is actually relatively easy to configure this feature. And how it is configured really depends on the details of how your server and organization are organized (that's why it's not preconfigured in a plug-n-play way). Personally, I love how SAS® runs on Unix servers, and when that's the case, I can leverage Unix Access Control Lists (ACLs) along with SAS® BI Server settings to configure SAS® and Enterprise Guide® to generate detailed SAS® log files in each user's own account on the server. And SAS® BI Server even enables me to configure this feature such that SAS® logs are automatically “rolled over,” for example, on a weekly basis.

No longer are the end-user analysts and statisticians left to fend for themselves or to keep their failed Enterprise Guide® application open on their desktop while they wait for a SAS® technical support “guru” to stop-by to see what might have been generated in their client application's SAS® log. By employing Unix server-specific features alongside SAS® BI Server configuration options, Zen Masters of the SAS® BI Server create a non-duality whereby SAS® logs are generated in real time on the SAS® BI Server, organized by end-user, and time-stamped. When receiving a phone call or an email asking for SAS® technical support for a problem Enterprise Guide® job, the SAS® “guru” can simply look at the user's log on the SAS® BI Server to locate error messages and diagnose the problem with ALL information available—from any location with access to the server, either onsite, from the building across the street, or any location in the world that supports at least a telnet (ssh) text-based login to the SAS® BI Server. This feature is readily available in the SAS® BI Server, but it's by no means “plug-n-play.” Creating it is unique to the specific environment at your organization.

Every SAS® BI deployment is a custom deployment.

## THE PATH TO A ZEN EXPERIENCE WITH YOUR SAS® BI SERVER

*Be master of your mind rather mastered by mind.  
-Zen Proverb*

Few software applications outside of the film and entertainment industry combine so deeply the hard-core aspects of enterprise-level information technology with the intellectual grace of statistical programming and data analytics. While it would seem, on the surface, that computer geeks and statisticians and analysts share a culture and a language,

nothing could be farther from the truth. Whereas the information technologist typically is process-oriented and functions in an environment that is, in many ways, “crowd control,” statisticians and analysts function in a more *ad hoc* environment, emphasizing procedure and path. It’s not in any way unusual for the two to be at odds with one another, not comprehending what often appears to be nonsensical language and behavior on the part of the other. The “anthropology of analytics” is a study of a very rich, colorful, and intense culture of technology, technique and tolerance.

At the heart of many successful SAS® BI Server deployments is a recognized effort at establishing either a formal or informal Business Intelligence Competency Center (BICC). The decision to deploy the SAS® BI Server isn’t merely a selection of a “tool” to apply to a problem. It’s the adoption of an entire technical and creative process that focuses resources with specialized skills toward those aspects of business intelligence for which they are best suited—in a manner that leverages their unique contributions across teams and departments. Examples of these specialized skills includes server and operating system management, database architecture, data management, application development, statistical analysis, business analytics, report design and generation, automation, and SAS® programming technical support as well as general computing technical support. And it’s perfectly alright to start small with the vision to build and expand efficiency over time.

To achieve a Zen Experience with your SAS® BI Server, organizations striving to get the most out of their software and infrastructure investment would leverage the skills and experience of those individuals who could contribute specialized SAS® BI Server managements skills as part of the BICC team. A SAS® BI Server architect and administrator should be able to take on responsibility for the architectural design of the SAS® BI system, design, develop, and constructs SAS® BI Server systems and their associated client solutions, consulting with stakeholders and end-users to refine application needs. S/he should be familiar with a variety of business intelligence concepts, practices, and procedures, relying on extensive experience and judgment to plan and accomplish goals.

An effective SAS® Enterprise BI System “guru” combines key aspects of IT infrastructure planning, enterprise database administration, operating system administration, educational psychology, end-user training, advanced SAS® macro coding techniques, and specialized metadata management. In short, s/he is an über-geek with exceptional people skills.

## ZEN MASTERS OF SAS® BI SERVER DEPLOYMENT

*It takes a wise man to learn from his mistakes, but an even wiser man to learn from others.*  
-Zen Proverb

*Best practices for implementing a BI strategy with SAS* by Mike Vanderlinden & Brian Varney

<http://www2.sas.com/proceedings/forum2007/245-2007.pdf>

*Best Practices for SAS® Business Intelligence Administrators: Using the Configuration Troubleshooter to Keep SAS® Solutions and SAS® BI Applications Running Smoothly* by Tanya Kalich

<http://support.sas.com/resources/papers/sgf2008/biadmin.pdf>

*Performance Tuning & Sizing Guide for SAS Users and Sun System Administrators [Updated Edition]* by William Kearns and Tom Keefer

<http://www.sas.com/partners/directory/sun/sugi28.pdf>

Business Intelligence Competency Center

<http://www.sas.com/consult/bicc.html>

## CONCLUSION

*Before enlightenment; chop wood, carry water. After enlightenment; chop wood, carry water.*  
-Zen Proverb

## ACKNOWLEDGMENTS

*Student says, "I am very discouraged. What should I do?" Master says, "encourage others."*  
-Zen Proverb

Credit goes to Darryl Holman, Ph.D., a truly priceless friend and ever-present personal influence on my earliest appreciation for a Zen approach to software engineering as well as to Unix systems administration. And with my apologies to Robert Pirsig, author of *Zen and the Art of Motorcycle Maintenance*, who's manuscript I do no credit with this missive.

## CONTACT INFORMATION

*The path of the enlightened one leaves no track- it is like the path of birds in the sky.*  
-Zen Proverb

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

Dave Steven  
Epistemata, LLC  
20126 Ballinger Way NE #75  
Shoreline, WA 98155  
dave@epistemata.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.