

# SAS<sup>®</sup> 9

## The First Enterprise-Class Business Intelligence Platform?

SAS describes SAS<sup>®</sup>9 as “the first Enterprise-Class Business Intelligence Platform”. What does that mean? Why is it important for your business? How close does SAS come to meeting the requirements implicit in such a description?

---

*Author: Philip Howard*

*Published: May 2004*

---



# Executive Summary

---

The business environment in which senior executives have to operate is becoming more and more demanding. On the one hand, there is ever increasing competitive pressure. On the other there is a growing requirement to meet corporate governance and compliance regulations. At the same time you need to drive down costs while simultaneously improving performance and productivity.

While squaring this circle involves the entire organisation, with all of its resources, of whatever type, an important component of this solution must involve software and, in particular, software that supports decision-making processes that are designed to improve operational efficiency and the alignment of the organisation with its strategic goals.

However, the provision of timely information needs to extend across the enterprise to all of its knowledge workers and beyond — something that offerings in this area have failed to do in the past. Further, such solutions must also address the governance and cost of ownership issues mentioned.

SAS®9 has been specifically designed to meet all of these needs. In this white paper we consider what would be required of a solution that would meet all of the needs outlined, and then we consider the SAS®9 offering and how it stands up to the criteria identified.

## Business Intelligence

It would be unfair to say that traditional business intelligence solutions have failed. Traditionally, business intelligence vendors have claimed that they could provide query and analysis capabilities that could be used across the organisation. However, that has not been true in practice. The reality is that most companies today have silos of separate solutions for separate user communities. They use reporting solutions for the mass of the workforce, analytics and OLAP (on-line analytic processing) tools for business analysts, data mining solutions and predictive analytics for specialists, and Microsoft Excel or some other spreadsheet product on a more or less help yourself basis. On top of this, it is often the case that different departments or divisions within a company may have selected different solutions in each of these areas (especially for reporting and OLAP).

While each of these solutions may be fine in its own right, their proliferation is a serious problem. It is not hard to spot the reasons for this. In principle, all of these solutions are complementary: they frequently rely on the same source data and the results are often used, or at least viewed, by people across the organisation (regardless of role). To put it simply, all of these business intelligence tools need to work together collaboratively.

There are, potentially, two ways to do this. The first is that you could, at least in theory, build a super-product that did everything: reporting, ad hoc queries, spread-

sheets, in-depth analysis, the lot. However, such a product would not be practical. Different users have different requirements and different capabilities and to build a product that had sufficient detail for business analysts and sufficient ease-of-use for line managers, all in a single package, would be a practical impossibility.

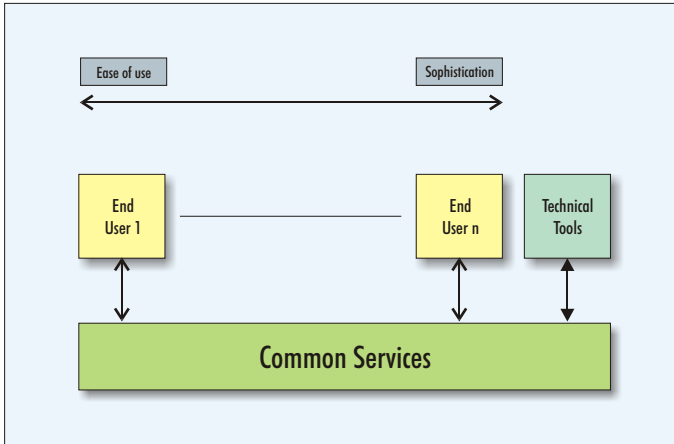


Figure 1: Technical applications sharing common services

The alternative is to provide a suite of interfaces, each designed with a particular user community in mind, but with a common look and feel, that is integrated under the covers. In this context, it is important to understand what we mean by “integrated under the covers”. To take a simple example; while there may be several user interfaces you only really need one reporting engine, you only need one scheduling facility, and it makes sense to have a common method for accessing data. In other words, you have a set of underlying facilities that serve as a common “platform” for the various front-end capabilities that are provided. It is also worth pointing out that it makes sense if this common platform also supports appropriate technical applications that also need access to these shared facilities. Such a solution would

look something like what is illustrated in Figure 1.

A secondary, but no less important, advantage of such an approach is that you can easily pass results from one end user environment to another, thus promoting collaborative working.

## Corporate Governance

Standards for corporate governance and compliance regulations abound. We do not intend to discuss any of these in detail. However, one of the key points in many of the legislative requirements is to be able to establish the veracity of published data. In particular, you need to be able to trace that data back to its source and see how it was subsequently aggregated and manipulated. This is referred to as Data Lineage. Further, you need to be able to demonstrate that you have appropriate mechanisms in place that can accurately trace the lineage of any piece of data.

The pre-requisite for providing data lineage is that you know where data is and what has been done to it. However, managing the data is not enough. You also need to be able to track and manage the metadata (data about data) that describes the original data. To put this simply, data is just a value: you might have a field that has “205” in it. That is an item of data. It is not until you know that that figure represents monthly sales of a particular product, say, that that data has any meaning.

Metadata actually goes beyond this simple example. It has implications for IT: for example, that “product sales” must be a numeric value, but it also has more profound consequences in business terms. For example, suppose that sales managers are asked to produce a report showing monthly sales and profits. The question is, how do you define sales? Does that mean orders or does it just mean orders that have been paid for? Further, what is profit? Does it include VAT or sales tax, or not? Does it include delivery charges or not?



Metadata can be used to define corporate standards for such things as what you mean by sales or profits and, once these definitions have been produced, they can be used and reused by managers and other personnel throughout the organisation. Note that this is desirable for its own sake, leaving aside any question of corporate governance.

How does this relate to data lineage? Because it is precisely such things as sales and profits that need to be traced, and the ability to do this depends upon having reliable metadata that can be used by the whole enterprise.

Data lineage therefore requires the ability to track and manage both data and the metadata that describes it. That, however, is by no means always a simple task. In particular, it becomes increasingly complex and difficult if you have multiple silos of products for different business intelligence needs. Data lineage is much easier to achieve when the environment is integrated and, in particular, when there is a single repository for all of your metadata. In other words, good corporate governance will be much easier to achieve when a single platform is used, just as it is with business intelligence.

This is not quite the end of the corporate governance story. In order to ensure data lineage is accurate, it is important that all data and metadata is fully controlled and that you have a complete audit trail of all changes that are made. This raises a particular problem with the use of Microsoft Excel and spreadsheet technology in general. Excel is not secure, is prone to errors (not in the software but in the formulae and calculations that are defined by the user), anybody who wants to can change data in a spreadsheet and, above all, there is no audit trail. While it is unlikely that anyone will stop using Excel anytime soon, it is therefore important that the use of Excel can be controlled within the business intelligence environment, or that comparable facilities are provided directly.

## Cost of Ownership

We have, by now, established the advantages of an integrated approach to business intelligence solutions, with a single platform and metadata repository serving a variety of front-end applications, tailored to the needs of different user communities. However, the benefits of such an approach go beyond ensuring that everyone is speaking the same language and meeting compliance requirements. In particular, there are two aspects of such an approach that should help to reduce the cost of ownership of such an integrated product set.

The first and most obvious advantage of integration from a cost perspective is that you are dealing with a single supplier rather than multiple vendors. This puts you in a better negotiating position, reduces administration and has other obvious benefits that we do not need to go into. However, that is not the end of the story.

Arguably, a more significant benefit deriving from an integrated approach is that the whole solution becomes much more manageable. This actually relates back to the metadata that we discussed above, as it is often this that entails the largest administrative overhead. If metadata is held in multiple places then ensuring that it remains synchronised can become a major overhead and, of course, this incurs costs. It is therefore beneficial to have an integrated platform with a single metadata repository purely from a cost point-of-view.



Moreover, the corollary is that the more products that work off this single repository the better. In practice this means that it will be advantageous if complementary technologies, such as those that extract data from source systems and load them into the data warehouse, and data quality tools of various sorts, are also supported by this same platform since this will, again, make the whole environment more manageable and reduce associated administrative costs. This is best illustrated by example. Suppose that you want to change one of your query applications, which accesses data from a variety of original sources. If it is easy to discover where all that original data is (which it will be in an integrated environment that provides impact analysis and similar facilities) then the maintenance of that application will be very much simpler, and therefore less costly, than in a discrete environment where such information is not available.

## Other considerations

While offering an integrated solution based upon a single metadata repository is a requirement for an enterprise-class business intelligence offering, it is not, in itself, sufficient. There are two other major requirements. The first of these is that the environment should provide sufficient facilities to support the enterprise's needs in terms of performance, scalability, resilience, and so forth. This should be self-evident.

Secondly, however extensive a solution of this type may be, it will not exist in isolation. In particular, there are a variety of complementary technologies that need to use the same metadata and business rules that will be used in the enterprise-class business intelligence platform. For example, data modelling, which is used to setup and define the data warehouse, does not typically fall within the area of competence of business intelligence vendors. Similarly, the ability (and it is an advantage, as we have seen) to load data from source systems together with facilities for ensuring the quality of that data, which will also need to use the same metadata, is not a pre-requisite for enterprise-class capability though it will be an advantage if this is provided by the same supplier. Another area where such openness is required is in front-end applications such as call centres, self-service web-based applications and others, where it is often the case that you need to access the business intelligence environment, for example to ensure consistent naming conventions (such as client names in a call centre) or to do access scoring mechanisms to support on-line credit risk checks.

In other words it is important that the environment provided by the business intelligence vendor supports open standards so that details can be shared with other environments and tools as easily as possible. In addition to standards for metadata, this particularly means support for open APIs (application programming interfaces) and Web Services.

## Summary

To conclude, an enterprise-class business intelligence solution has two key characteristics from which all other requirements flow. First, it must offer a wide range of different capabilities that have been designed to support the needs of different user communities within the organisation; and second, all of these facilities must be integrated through a single metadata repository that is open to outside access. Scalability,



resilience and performance are also pre-requisites but these are hardly unique to a business intelligence environment.

In the following section we will describe SAS®9 and see how it meets the criteria we have discussed for an enterprise-class platform.

# Product description

## Introduction

Although technically a part of the business intelligence community, SAS has always been best known for its statistical and data mining solutions on the one hand, and its industry-specific analytical solutions on the other. While it has offered a variety of business intelligence products for some time, there has remained a perception that SAS was not really in the same space as companies such as Business Objects and Cognos. With SAS®9 (which is actually in version 9.1) SAS has set out its stall to compete directly with these vendors.

However, SAS®9 is more than just a business intelligence platform. It is, in fact, the platform for the entire SAS solution set, so that all of its analytic applications, its data and text mining products, its clickstream analysis and everything else that it does, including business intelligence, will be based upon the SAS®9 platform.

It should be appreciated that developing SAS®9 has not been a trivial task. It was announced as long ago as 2002 and, no doubt, it was being planned long before that, so a lot of man years have gone into its development. In particular, prior to this release, there was no single metadata repository that underpinned the entire product set. If you had any diverse set of SAS products you would have multiple repositories and, while you could pass information from one to another, there were the inevitable complexities of having to manage these repositories and the movement of metadata between them. By developing an integrated, single platform solution this complexity has simply been removed, though that does not mean that all complexity has been eliminated, because there remain issues in interoperating with other environments, though this is a failure in the current standards that are available rather than anything specific to SAS.

The fact that SAS®9 provides an integrated platform will be a significant boon for existing and new users of traditional SAS products. However, in this white paper we are interested in SAS®9 within the context of business intelligence. The truth is that without the single platform approach provided by SAS®9, the company could not legitimately offer the sort of comprehensive, enterprise-wide business intelligence solution that it is offering.

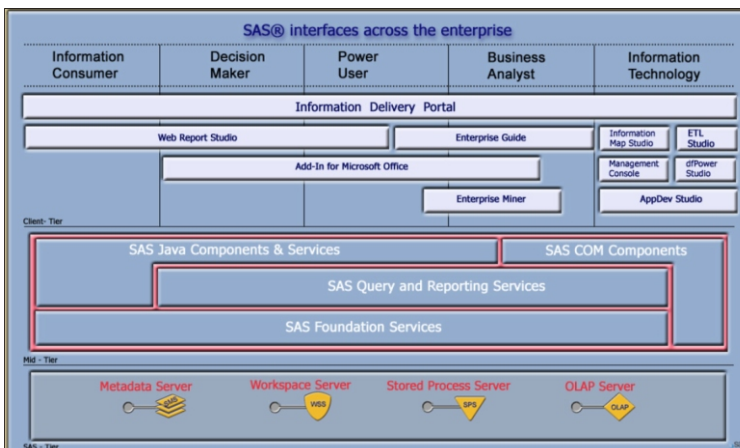


Figure 2: SAS®9 architecture

## Architecture

Figure 2 above illustrates the architecture of SAS®9 within the context of business intelligence. However, for the purposes of this white paper, we are not going to discuss all of the elements shown here. Nevertheless, this diagram is useful because it illustrates the different user communities that SAS is addressing, and the products that



are applicable to each of these sectors, notably SAS Web Report Studio, the SAS Add-in for Microsoft Office, and SAS Enterprise Guide®, which we will discuss below.

It is also useful in that it identifies the fact that SAS®9 uses different servers for different purposes, for example a separate OLAP Server and Metadata Server. Further, it indicates the three tier architecture that has been employed although not the fact that actually an n-tier environment is envisaged. For example, the SAS Information

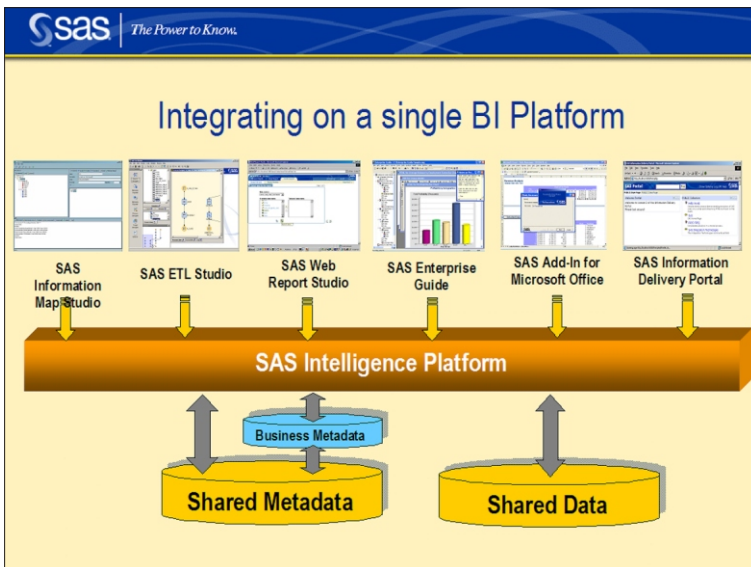


Figure 3: A single BI platform

Delivery Portal is J2EE compliant and it will be resident on an application server, with end users being browser-based. As far as the servers themselves are concerned, these may be deployed on mainframe (z/OS), Windows, UNIX (HP-UX, AIX and Solaris) and Linux platforms, and support a broad range of standards running from J2EE and COM through LDAP directories to SOAP and UDDI, and from messaging environments such as WebSphere MQ and MSMQ to OLE DB and OLE DB for OLAP.

Referring back to Figure 2, we will not be discussing the mid-tier components and nor will we be discussing SAS Enterprise Miner®, which is the company's data mining and predictive analytics tool. In practice, what we will be discussing in more detail is better illustrated in Figure 3.

As can be seen, the various front-end products (of which SAS Information Map Studio and SAS ETL Studio are designed for use by the IT department rather than end users) all run against the same metadata repository and access the same data, through a common set of servers. In more general terms, we will discuss SAS®9's back-end capabilities later and start by considering the various tools provided for the different front-end user communities.

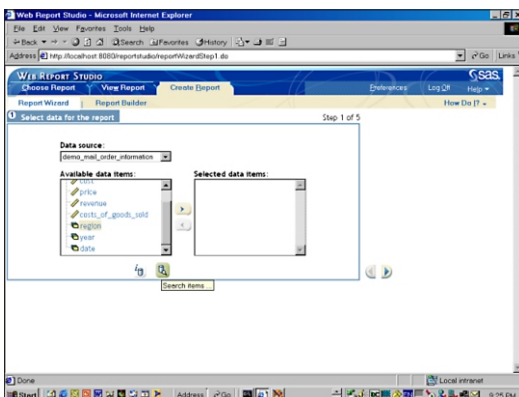


Figure 4: Screenshot of Web Report Studio

## SAS® Web Report Studio

SAS Web Report Studio does what it says on the tin. That is, it is a browser-based product for viewing and creating reports. Some people will use it simply to view reports while others, depending on their role, will be allowed to create reports that can be viewed by other users. In the latter case, a wizard is provided to help you create reports while more experienced users can employ the Report Builder. Note that because of the business metadata that is stored by SAS®9 (see later), all the descriptions of data sources and the data items that you can report on are expressed in business terminology rather than technical jargon.

## SAS® Add-in for Microsoft Office

The SAS Add-in for Microsoft Office offers a number of facilities. For example, it allows you to include a report from SAS Web Report Studio within a Microsoft Word document, or a business intelligence graphic or chart from SAS Enterprise Guide



within a PowerPoint presentation. Moreover, since the graphic may be dynamically updated, this has the significant advantage that you can be presenting with dynamically updated data.

The way that this collaboration works (and, indeed, it applies across all the elements of SAS®9) is that the product makes use of stored processes. These are a way of encapsulating business rules. Such a business rule might simply lead to a derived data field or it might result in a report or a chart or whatever is appropriate. In any case, the rule (stored process) is held within the SAS Metadata Repository so that it can be reused by any participating tool. So, for example, you could use SAS Enterprise Guide to publish the results of a data mining exercise within SAS Enterprise Miner, or you could combine the latter with ETL capabilities. In effect, you can move up and down the levels of sophistication within the product set, as required.

Note that what is actually happening here is that the presentation of information has been separated from the business logic: the stored process defines the latter, while the user interface defines how the information will be presented. This has the advantage that can you change one without impacting on the other.

In the case of the SAS Add-in for Microsoft Office there is a mechanism for directly passing this logic defined in a stored process to the relevant part of Office. For external applications, stored processes can also be exposed as Web services. Of course, the most common use of the integration provided by SAS will be with Microsoft Excel, whereby SAS provides automated facilities to export data generated using any of its other products into Microsoft Excel, where the data can be manipulated further. Note, however, that once that data has been exported, it is within the Excel domain and outside of the SAS environment.

## SAS® Enterprise Guide®

SAS Enterprise Guide is the company's product that fits most clearly into the con-

ventional business intelligence and OLAP space inhabited by the likes of Business Objects. As indicated in Figure 2, it uses (or may use-see later) its own OLAP Server and it has exactly the sort of slice-and-dice, pivot, drill down, and other capabilities that you would expect. As illustrated in Figure 5, you can output reports and charts in a variety of formats. In particular, given that SAS has always been well-known for the excellence of both its statistics and the graphical representation thereof, it is not surprising that SAS Enterprise Guide is also strong in these areas. Note the ability to append notes to charts and reports, which will be useful in collaborative environments.

In more technical terms, SAS supports both multi-dimensional and hybrid data stores and there is native support for OLE DB for OLAP, while the query language used is MDX (multi-dimensional expressions).

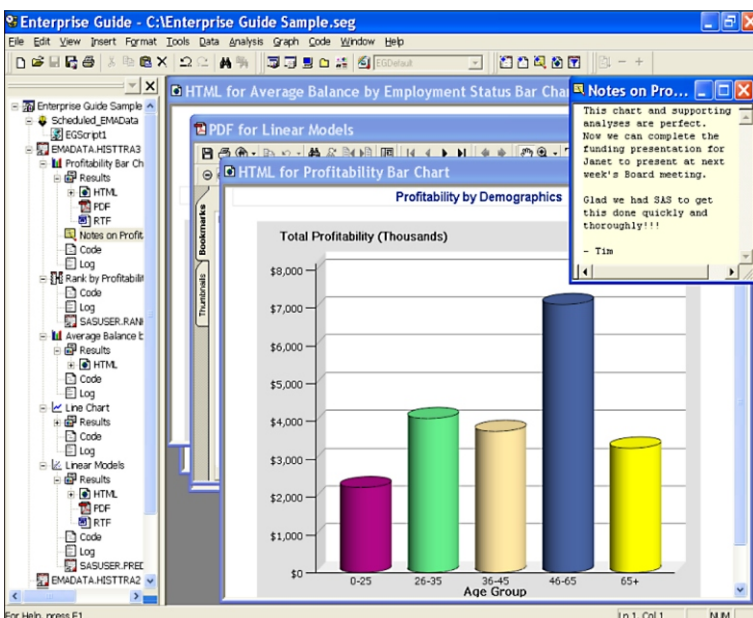


Figure 5: Screenshot of Enterprise Guide®



This has important consequences. Because SAS' OLAP facilities are standards-based, this means that you can use SAS Enterprise Guide with a third party OLAP server, if you prefer. Alternatively, you can also use third party front-end tools in conjunction with SAS OLAP Server.

Other facilities within the product include a cube-building tool called SAS® OLAP Cube Studio, which works closely in conjunction with SAS ETL Studio (see later) so that you can build cubes directly from cleansed data. The product supports time dimensions, ragged and unbalanced hierarchies, and parallel drill hierarchies that allow you to analyse along different drill paths within a single dimension.

## SAS® Information Delivery Portal

This is a fairly standard portal that acts as a delivery mechanism for users of all sorts of output. This may include the reports, graphics, and so forth that are produced from the front-end products just discussed; but it can also include dashboards, balanced scorecards and details derived from other SAS products, as well as information from other sources. Included facilities are a search engine, navigator portlets that allow you to browse the repository (see below) directly from within the portal, role-based security, notification portlets that surface alerts and can be used in conjunction with KPIs (key performance indicators) to notify relevant parties when a KPI in a dashboard goes out of range, encryption capability, single sign-on, and a development environment for constructing portlets called SAS AppDev Studio™.

As one would expect, there are full personalisation and customisation capabilities and the portal has already been internationalised with support for a number of different language versions, including English, German, Italian, French, Spanish, Chinese and Hebrew.

The SAS Information Delivery Portal uses a conventional portlet-based approach and you can, of course, deliver results to a third party portal if you prefer, or you can implement the SAS portal as a part of a wider enterprise portal environment.

## SAS® ETL<sup>Q</sup>

In Figure 3, there are two major additional products that are shown as a part of the SAS®9 architecture: the SAS Information Map Studio and SAS ETL Studio. However, the former works in conjunction with the SAS Metadata Repository and is discussed below, while the latter is actually a part of a larger product category in its own right. SAS markets this as SAS ETL<sup>Q</sup>, where this is the company's integrated data integration platform, consisting of ETL (extract, transform and load), data profiling and data cleansing.

SAS ETL Studio is a template-based development environment that supports the design of data transformation and mapping processes within the ETL environment. It is packaged in two ways, either as part of SAS Enterprise ETL Server, or SAS ETL Server depending upon the requirements and scale of the organisation concerned.

While you can use SAS ETL Studio in conjunction with third party data quality tools, the integration between these will obviously not be as tight. For example, in SAS



ETL Studio there are built-in profiling and cleansing icons that you can drag and drop onto your palette while developing transformation workflows; if you are using third party products these will not be available and you will have to code such routines. Further, using third party tools will immediately mean having a second metadata repository and a third if profiling and cleansing are sourced from different suppliers. This will make data lineage more complex and increase management overheads. In addition, there are significant performance advantages to be gained from an integrated environment. For example, SAS simply streams data to its cleansing facilities as it is processed; in non-integrated environments you typically have to write the data to a file, which is subsequently read by the data cleansing product. For large files this is a process that can take a matter of hours and could severely impact on batch windows and performance overload.

As one might expect, there are a significant number of connectors provided by SAS to support a variety of data and other sources, including specialised facilities for extracting data from various application environments and support for message queuing. It should be noted that these connectors are reused across the SAS®9 product set, thus the same facilities could be used in the SAS Information Delivery Portal for accessing data from an SAP R/3 system, for example. Further, these connectors can be re-deployed by any relevant product in a live (federated) environment. Thus, for example, one of the business intelligence tools can use a message queuing connector to directly pass the results of a query to another environment via WebSphere MQ or a similar product. There is also an ODBC bridge provided by SAS to support integration with other external environments.

## Metadata Repository

The development of this part of SAS®9 represents a major achievement. In the past, different SAS products had separate repositories and, though they could talk to one another, this is nowhere near as convenient as sharing a single repository. The multiple repository scenario is the position that most other business intelligence vendors are still in today.

Note that one consequence of this approach is that you can use the various tools discussed above against the repository for browsing, searching or other purposes. For example, you might want to monitor the number of errors that were being detected by your data quality procedures (assuming you are using the data quality in SAS ETL Studio) to support quality assurance programmes designed to reduce such error rates. Or you could use them to see how popular particular reports were and which were viewed most frequently.

The SAS Metadata Server, which provides repository capabilities for SAS®9, is open and supports the OMG's (Object Management Group's) Common Warehouse Metamodel (CWM) standard. In particular, it supports the use of XML for exchanging metadata with other environments, and this is where the SASXML Mapper comes in, since it enables interrogation and browsing of the repository by means of XML, together with mapping capabilities that enable interoperability.

Apart from CWM, SAS also has a partnership with Meta Integration Technology Inc (MITI), which provides meta-bridges that are integrated with SAS®9 to allow you to



exchange metadata that is stored in a proprietary format in various applications and third party environments.

The SAS Metadata Server, which stores both technical and business metadata, provides check-in and check-out, versioning, impact analysis, access control, and a complete audit history of who has done what and when. There are also import and export wizards for occasions when you want to import or export metadata directly into or out of the environment without going through SAS ETL Studio or via SAS Information Map Studio.

Finally, the SAS environment offers an open API that can be addressed via a number of environments such as XML, Java and COM/DCOM. This is important because it enables front-end applications to directly address the SAS®9 environment. For example, suppose that you have standardised on IBM as a business description of that company, as opposed to International Business Machines. If an employee at IBM phones into a call centre you want to ensure that the call centre records his or her employer as IBM. These sorts of standardisation rules will normally be determined using the data quality software in ETL<sup>Q</sup> and then stored for reuse within the repository. The Open API means that the call centre application can access this information directly. Similarly, while we are not discussing the use of Enterprise Miner in this white paper, things like credit risk scoring or fraud algorithms that you can develop using Enterprise Miner can be stored and, again, appropriate front-end applications can access these routines directly through the open API. Alternatively, these resources can be exposed as Web services and this mechanism can be used for integration with the front-end environment as an alternative to using an API. Note that standard names, risk algorithms and so forth are not stored in the Metadata Repository themselves, but the Metadata Repository knows where they are held, thus providing a common addressing facility across the product set.

## Manageability

Clearly, much of the manageability in SAS®9 is derived from the fact that it is a single environment. However, that is supported by the SAS Management Console, which acts as a single point of control for the whole environment through which all the different lower level managers run. These all do more or less what their names suggests and include the Metadata Manager, which controls all of the metadata servers (you can have a federated metadata environment), their repositories, resource templates, change management and import/export; the Server Manager, which is used to define and configure SAS servers and connections; and the User and Authorization Managers. In a multi-machine cluster or grid environment the SAS Metadata Server(s) act(s) to bind these together and the SAS Management Console acts as the window into the Metadata Server(s)

## Scalability, Resilience and Performance

The operative word when it comes to discussing the sort of scalability, and other features that all large-scale systems require, is flexibility. For example, SAS ETL<sup>Q</sup> does not enforce a strictly E, T and then L process upon the user. If it makes more sense from a performance perspective to do some transformations (join operations say) on the source system before extracting the data, then you can do that. Similarly, you



could load all the data onto the target and then transform it: that is what we mean by flexibility.

Another aspect of flexibility which we have already mentioned is the range of platforms that are supported by SAS. In the business intelligence world its support for mainframe environments is a key differentiator. Moreover, it offers specialised servers for different functions. For example, the SAS® OLAP Server, which is multi-threaded (as are all SAS servers), has been significantly re-designed in this release to provide improved performance. Other standard features include load balancing, failover, independent parallelism, pipeline parallelism and so forth. Perhaps most significantly, in SAS®9 you can distribute either an application or the platform itself across a grid or cluster, so that you can make optimal use of existing resources.

In addition to supporting mainframe environments, SAS also supports mixed environments (another differentiator) in which you can have both mainframe and other systems co-operating with one another. In this sort of environment SAS uses a costing algorithm that understands that mainframe operations are more costly than those on other platforms and it therefore load balances asymmetrically, depending on the results of the algorithm.



# Conclusions

There are really two sets of points to consider: how well does SAS®9, as a business intelligence platform, answer the sorts of questions we posed at the beginning of this paper? And, how well it does it stack up to the competition in answering those questions? After all, SAS might meet the criteria for the first but so might lots of other vendors, so the initial question needs to be put in context.

First, then, how does SAS®9 compare to the requirements identified?

	Requirement	SAS®9
1	Business Intelligence – integrated suite of user interfaces and tools that span the enterprise	Yes
2	Corporate Governance – integrated metadata environment to provide data lineage	Yes
3	Corporate Governance – control of Microsoft Excel environment	No
4	Cost of Ownership – single vendor	Yes
5	Cost of Ownership – improved manageability through single metadata environment	Yes
6	Other Considerations – performance, scalability and resilience	Yes
7	Other Considerations – interoperability	Yes

Table 1: Does SAS compare favourably to requirements?

As can be seen, SAS®9 fulfils all of our requirements for an enterprise-class integrated business intelligence platform with the exception that it does not provide control of the Microsoft Excel environment, which is a feature that we would certainly like to see the company introduce. In the meantime, the best we can suggest is that users make as much use of stored processes as possible when working within an Excel environment, and limit the re-definition of data within Excel as much as is feasible.

As far as its competitive position is concerned, the following table shows the current state of the market against which SAS®9 has to be compared. The numbers shown relate to the same requirements as in Table 1.

From Table 2 it should be clear that, while we may have reservations about the exploitation of Excel, in all other respects SAS®9 is either up to or ahead of the market. However, we should note the proviso that we are not comparing the individual features of the SAS products here, only the general environment. There may well be individual features of competitive products that we would like when compared to SAS (and the reverse is also true, of course) but we are not considering the competitive landscape at this level of detail.

What we can say is that there is no other suite of products in the market that has the breadth of capability that SAS®9 can provide with the level of integration that underpins it. Put that together with the scalability and other features that the platform pro-



How the market fares	
1	Most vendors are struggling towards the position that SAS has reached. We know of only one other fully integrated set of products based on a single platform but this has limitations with regard to its OLAP tool and, in any case, it is much narrower product set overall.
2	Most suppliers can provide this, even where separate metadata repositories are required by tracking the movement of metadata, though this is obviously more complex.
3	A number of vendors can control the Excel environment or have developed their own Excel look-and-feel products. However, there remain a substantial number of suppliers that cannot control Excel capability.
4	Nobody has as wide a range of business intelligence solutions as SAS.
5	Most other suppliers are trying to reach the same position as SAS. Many of them are years away from an integrated solution.
6	No other company in the business intelligence space can match SAS for across-the-board scalability.
7	SAS is comparable to a number of other companies in terms of interoperability and we do not regard this as a differentiator.

Table 2: SAS compared to the market

vides and we concur with the company that SAS has produced an enterprise-class business intelligence platform. We would also agree that it is the first such product in this class and that this should give the company a substantial advantage as it can now move beyond its traditional high-end analytics stronghold to address the general-purpose business intelligence market for the first time.

# Copyright & Disclaimer

---

This document is subject to copyright. No part of this publication may be reproduced by any method whatsoever without the prior consent of Bloor Research

Due to the nature of this material, numerous hardware and software products have been mentioned by name. In the majority, if not all, of the cases, these product names are claimed as trademarks by the companies that manufacture the products. It is not Bloor Research's intent to claim these names or trademarks as our own.

Whilst every care has been taken in the preparation of this document to ensure that the information is correct, the publishers cannot accept responsibility for any errors or omissions.



Suite 6, Challenge House, Sherwood Drive,  
Bletchley, Milton Keynes, MK3 6DP, United Kingdom

Tel: +44 (0) 1908 625100 - Fax: +44 (0) 1908 625124  
Web: [www.bloor-research.com](http://www.bloor-research.com) - email: [info@bloor-research.com](mailto:info@bloor-research.com)