

Paper 483-2013

Integrating SAS into your operational environment: SOA as “a means to an end”

Saravana R Chandran, SAS Institute Inc., Cary, NC

ABSTRACT

The impacts of business analytic models have proven value for enterprises. Many SAS customers have highly valuable analytical assets ranging from Analytical Models to Analytical Services specific to a domain. We have seen a significant stream of requests from customers for assistance in taking the next step and deploying these models into their primary operational business applications. SOA (Service Oriented Architecture) is an architectural style designed to enable Flexibility, Reusability and Interoperability and provides one of the primary means for integrating SAS with your operational application environment. The paper walks through the Integration, Runtime Environment, Governance and Best Practices all in the context of SOA and SAS Business Analytics.

INTRODUCTION

Service Oriented Architecture's direct business benefits are reduced costs, reduced redundancy, better consistency, and improved efficiency through its reusability and flexibility. Integrating SOA based Analytical services into operational environments requires adapting to certain best practices as part of planning, designing, implementing and maintaining.

The paper describes the how to leverage SOA and the Analytics Life Cycle together and provides examples of how you can use SAS to build and use both synchronous and asynchronous services end points.

Service Oriented Architecture (SOA) is an architectural style designed to enable flexibility, reusability and interoperability in which the fundamental unit of design is reusable services. The OASIS SOA Reference Model defines SOA as:

“Service Oriented Architecture (SOA) is a paradigm for organizing and using distributed capabilities that may be under the control of different ownership domains. It is natural in such a context to think of one person's needs being met by capabilities offered by someone else; or, in the world of distributed computing, one computer agent's requirements being met by a computer agent belonging to a different owner. There is not necessarily a one-to-one correlation between needs and capabilities; the granularity of needs and capabilities vary from fundamental to complex, and any given need may require the combining of numerous capabilities while any single capability may address more than one need. The perceived value of SOA is that it provides a powerful framework for matching needs and capabilities and for combining capabilities to address those needs.”

Most enterprises have analytical assets ranging from analytical models to analytical services that they use in their decision making processes. SOA and its principles bring significant value for to the enterprise with respect to automating and integrating enterprise analytical assets in operational applications. Integrating analytical processes into the operational environment is a critical need and differentiator for today's enterprise. Fortunately, the Analytics Life Cycle and SOA life cycle complement each other and thus can provide a significant time savings and add flexibility into integrating analytics into a company's operational processes.

THE ANALYTICS LIFE CYCLE IN ENTERPRISE

Enterprises that organize business processes to align with the Analytics Life Cycle improve their business agility because they have flexibility in adapting to change. Although each enterprise organization approaches analytics differently, at the core the Analytics Life Cycle is the same. Figure 1 below shows various phases of the Analytics Life Cycle.

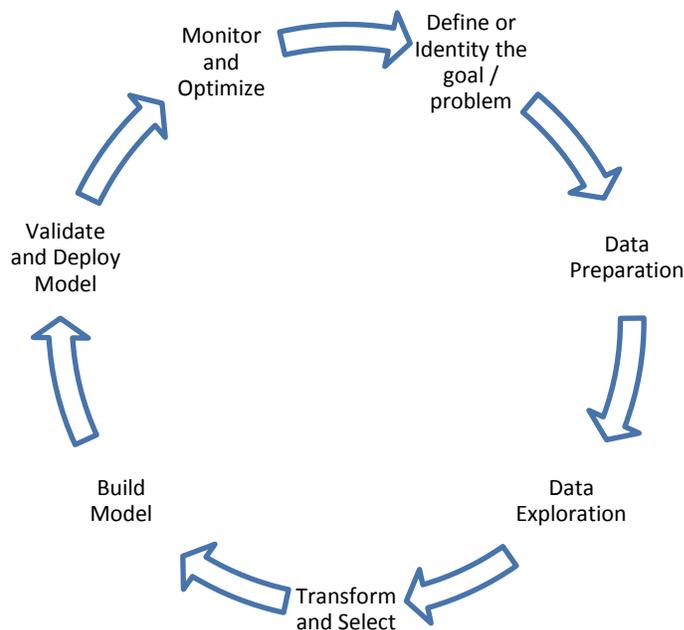


Figure 1 - The Analytics Life Cycle

Inconsistent use and management of analytical models creates challenges for many organizations, and increases the time and resources spent on putting models into production. Spending large amounts of time moving models through the dev-test-prod cycles can result in models that are “stale” by the time they reach production – or in models that never get deployed. Building a model requires many iterative runs involving selection of variables, testing and validation to model which ones have the most impact or greatest likelihood of indicating an outcome. These models are then compared to each other over time to monitor the effectiveness of the model and to update or retire models when they are no longer meeting requirements. Managing these models requires an understanding of the life-cycle stage of the model and keeping track of who touches the model and when. The analytical life cycle and SOA Service Orientation life cycle have many similarities and complement each other quite well. Model development and deployment are critical building blocks when developing SOA based Analytical Services for use across the enterprise.

WHY DELIVER ANALYTICS AS A SERVICE?

Certain types of analytical services need to be integrated into business processes across the enterprise in order to have impact for the business in right time and right place. In many enterprises, business analytics leverage batch processing and off-line data analysis. Transactional systems such as sales, marketing, CRM, finance, and procurement capture and produce transactional data which is then pulled into data warehouse and data mart projects. The analyst works on the offline data to build models, and perform comparison and validation analysis before making recommendations to the business. This process takes place outside of the operational application execution process and changes or improvements in models often encounter a significant lag time before being reflected in the operational applications. However, some types of operational applications can gain significant value by consuming near real time analytical processes “in-line” and surfacing the most recent model output directly into operational processes.

SOA based analytical services enable integration with various business user business intelligence and analytic applications such as dashboards, portals, mobile applications and to integrated event system based applications. Regardless of the need for real time or batch oriented use, the SOA based analytical services are important assets to help make fact based decision making a part of business processes execution in operational applications.

SAS ANALYTICS AND SOA FRAMEWORK

SAS Analytics and the SOA framework shown below are divided into three phases related to the Analytics Life Cycle – Source, Discover and Share. The framework identifies four key business activities associated with each phase:

functionality, value, stakeholders and technology choices. In this paper we focus on the “Share” phase and describe how service oriented architecture is used to share the insight produced by business analytics. During the Share phase you will be building analytical services based on models developed in the Source and Discover phases. The task of Share in this case is integration into the operational environment. Service orientation is the technology choice for implementing the SOA based analytical services.

Not all analytical models and assets end up as SOA based analytical services. When the results or output of analytical models can be interpreted programmatically then it becomes a candidate to be “shared” as analytical service. In general, supervised models are candidates for SOA based analytical services. Most often SOA based analytical services that enable decision making earlier in the business process tend to have more value in integrating into the operational application environment. SOA based analytical services are typically created to automate and standardize a critical part of a business process.

| Business Analytics Activity | Source | Discover | Share |
|-----------------------------|--|--|---|
| Functionality | <ul style="list-style-type: none"> ▪ Data capture ▪ Data cleansing ▪ Data manipulations ▪ Master data management | <ul style="list-style-type: none"> ▪ Query ▪ Statistical & analytical models ▪ Scoring ▪ Visualization | <ul style="list-style-type: none"> ▪ Publishing ▪ Interaction ▪ Integration ▪ Reusability |
| Value | Provide consistent data across the enterprise <i>“Know what you know”</i> | Uncover opportunities <i>“Realize what you didn’t realize”</i> | Distribute actionable intelligence <i>“Deliver efficiently”</i> |
| Stakeholders | <ul style="list-style-type: none"> ▪ IT Data Management ▪ Operational systems | <ul style="list-style-type: none"> ▪ Business Decision Makers ▪ Decision support systems | <ul style="list-style-type: none"> ▪ Business interfaces ▪ Business processes |
| Technology | <ul style="list-style-type: none"> ▪ Data access tools ▪ Data quality algorithms ▪ Custom transformations ▪ Job flow management and scheduling | <ul style="list-style-type: none"> ▪ Modeling language flexibility ▪ Analytical modeling ▪ Custom development ▪ Interactive diagrams | <ul style="list-style-type: none"> ▪ Services orientation |

Table 1- SAS Analytics and SOA Framework

SERVICE ORIENTATION

The service orientation life cycle has four critical phases - Planning, Modeling, Implementation and Management.

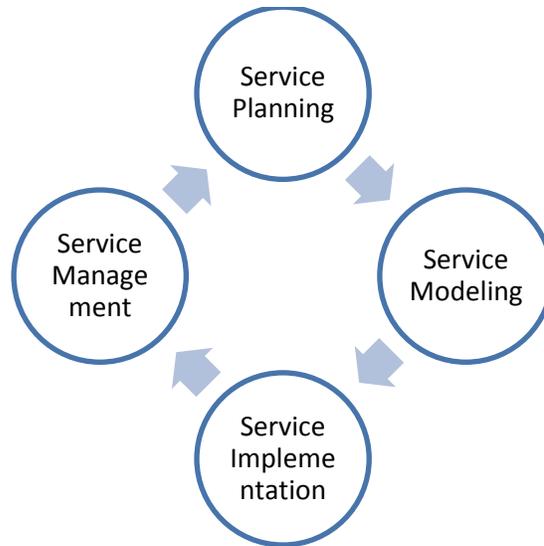


Figure 2 - Service Orientation Life Cycle

When you want to create an analytical service based on a model which has already gone through the Analytics Life Cycle for model development, validation and deployment, then you can simply overlap the service planning and modeling with the model development life cycle. The figure below shows the alignment and integration between the life cycles for an analytic model and for SOA.

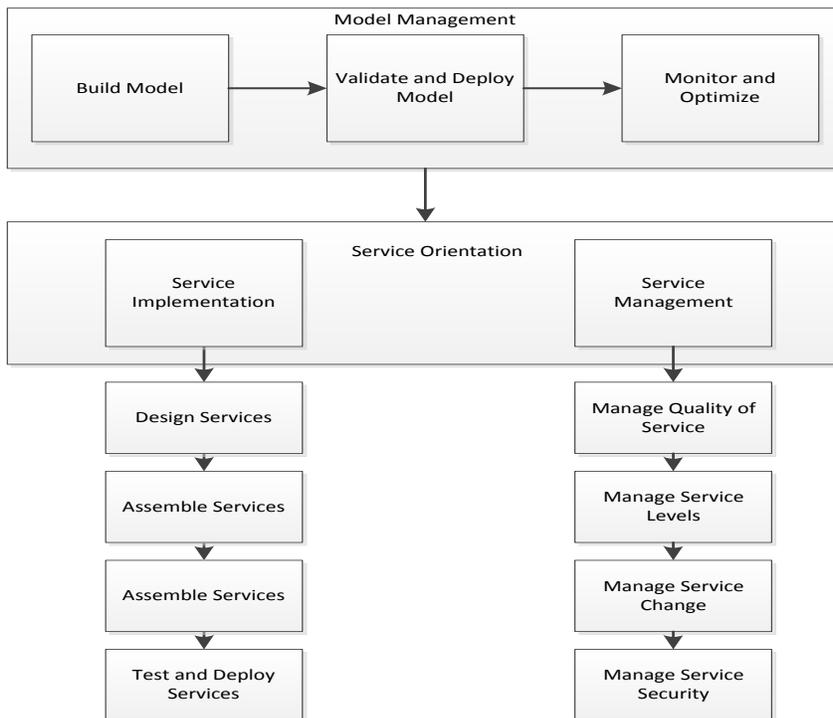


Figure 3 - Integrated SOA Based Analytics Service Life Cycle

INTEGRATING ANALYTICAL SERVICE

There are several approaches for integrating analytical services into the operational environment. From an integration perspective, SAS supports both synchronous and asynchronous integration of services. You would use

synchronous integration when your integration pattern is request and reply. And when your integration pattern is fire and forget, then you would typically use an asynchronous integration model. Business Process requirements determine the integration pattern.

A common approach for integrating SAS analytics into an operational environment is to package SAS code as a stored process that can then be exposed as a web service. A business process or application can call that web service to activate the SAS code and retrieve results. Another common approach is to build scoring service based on SAS tools such as SAS Enterprise Miner®, SAS/STAT® software, and Base SAS. These tools can be used to create scoring code in many different formats, such as SAS DATA step, C, C++, Java, and PMML (Predictive Model Markup Language). In some case PMML code is exported from SAS Model Manager and deployed into a non-SAS environment such as a database. Other options include pushing a scoring algorithm or other analytical model to the database using approaches like the SAS Scoring Accelerator in conjunction with SAS Access Engines. The format of the scoring model depends on the needs of scoring execution environments. With SOA integration, however, the scoring program interface (SPI) remains the same regardless of the underlying target scoring execution environment. Using the standard SPI interface, you can integrate your business process with established and standardized web services.

SYNCHRONOUS END-POINT

The diagram below shows the overall development and integration approach when you want to use a synchronous end-point. During service planning phases you might identify the need for an analytical service such scoring service.. To satisfy that service specification, you might identify existing reusable SAS programs or write new SAS programs that perform the desired analysis. You would then package that SAS code as a stored process and expose it as a web service.

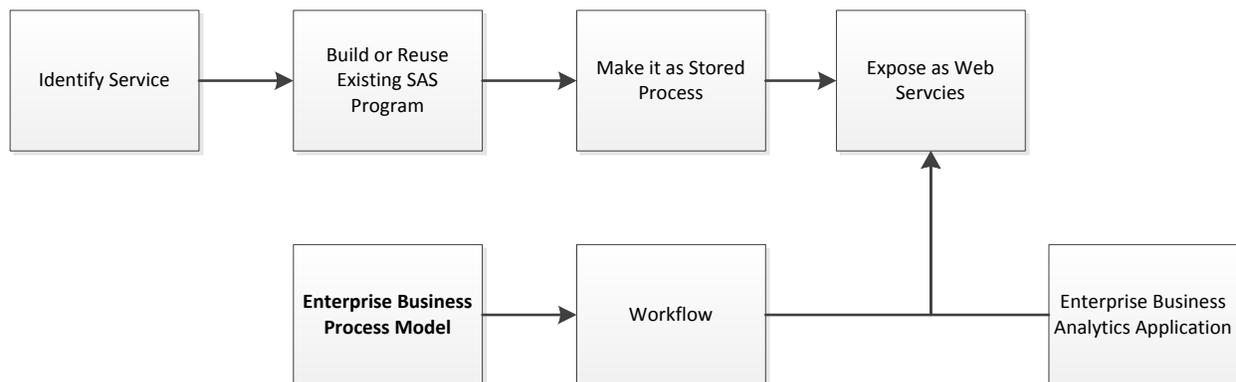


Figure 4 - SOA Analytical Service as synchronous endpoint

Many SAS applications can be used to create Stored Processes, including:

- SAS® Management Console to register SAS code
- SAS® Data Integration Studio (previously known as SAS® ETL Studio)
- SAS® Enterprise Guide®
- SAS® Forecast Server

ASYNCHRONOUS END-POINT

SAS supports asynchronous end-point delivery of services through a messaging interface. For example, to deliver an analytical scoring service, you would first build, test and validate a scoring model using tools such as SAS Enterprise Miner or Base SAS DATA Step code. Using SAS Integration Technologies, you could deploy the model with a message interface and that enables external service integration.

The diagram below shows the overall flow of the development and integration of a service accessed via an asynchronous end point.

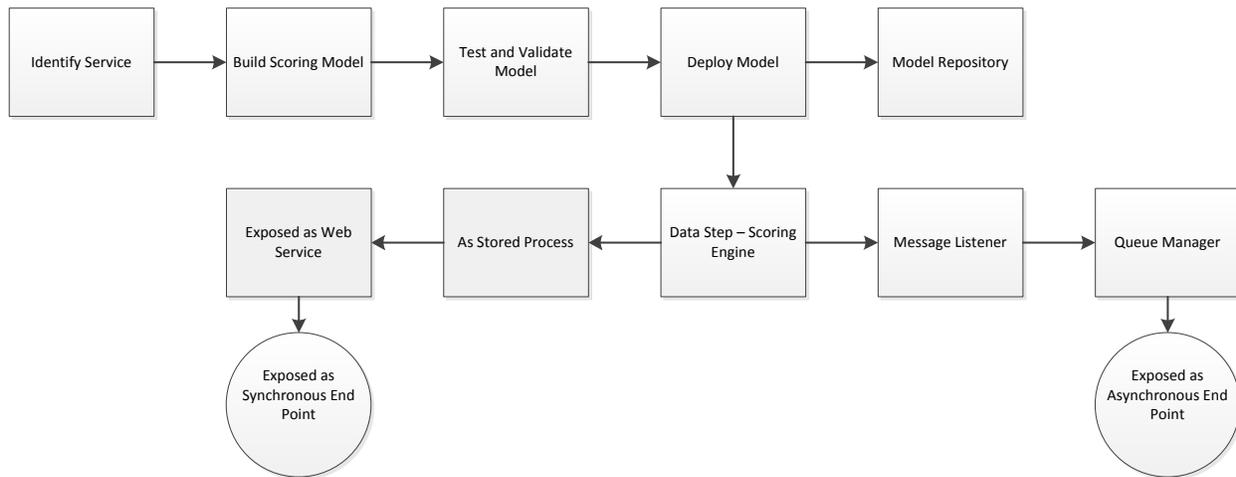


Figure 5 - SOA Analytical Service as asynchronous endpoint

RUNTIME ENVIRONMENT

SOA-based analytical services run in a SAS environment that includes a metadata server and a SAS middle tier. The SAS middle tier includes the Web Infrastructure Platform, which enables integration via a services registry used for querying and binding services – both those provided by SAS and those created to meet a specific enterprise need.

The tiers in the SAS topology are consistent across SAS applications and solutions.

- **DATA** - The data tier represents the data sources of both SAS and Non-SAS.
- **Compute** - The compute tier runs one or more SAS Servers. When running SOA Services, the SAS Workspace Server, SAS Stored Process Server, SAS Metadata Server and Object Spawner are used heavily during the runtime.
- **Middle** - The SAS middle tier includes the web applications that are deployed into a Web Application Server, SAS Remote Services, SAS Web Infrastructure Platform (WIP), and SAS BI Web Services.

Depending on the service need that was identified during your initial business analysis, you might use various SAS clients to create and deploy an analytical service. In the above stated example (Asynchronous end point) you would end up using SAS Model Manager and SAS Enterprise Miner as driver of service identification and modeling.

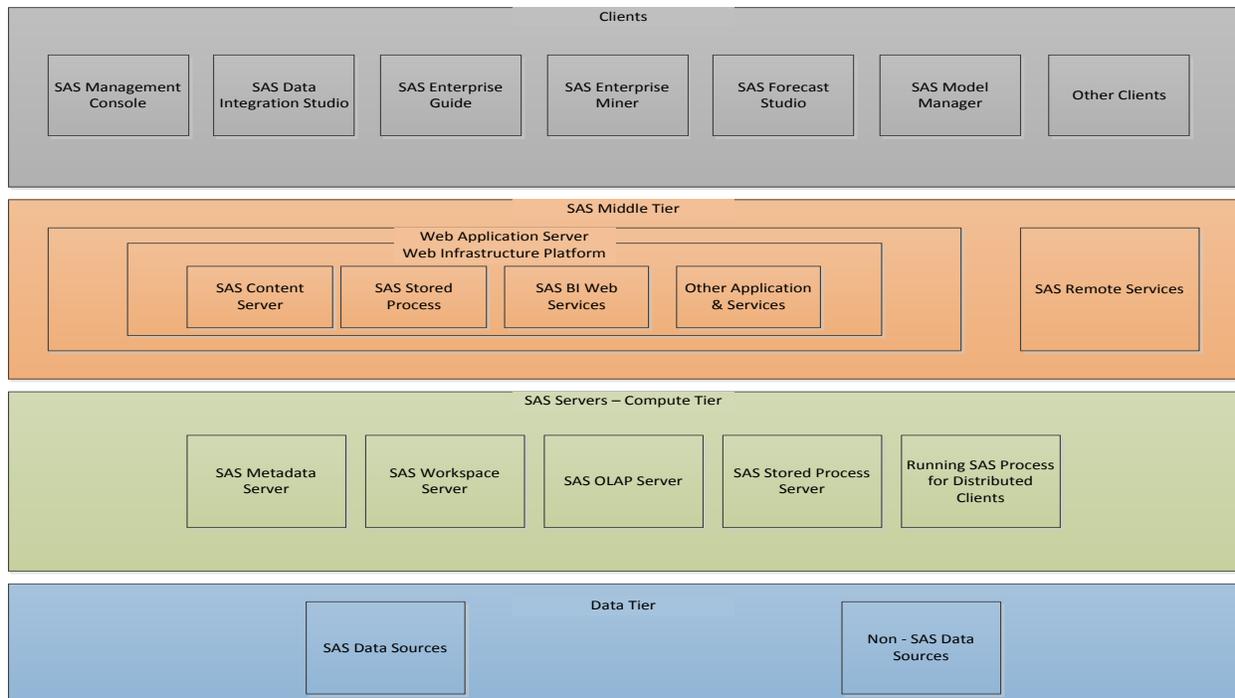


Figure 6 - SAS SOA Runtime Environment

BEST PRACTICES

Integrating SOA based Analytical services into operational environment requires adapting to certain best practices as part of planning, designing, implementing and maintaining. The following list highlights important best practices.

Meaningful to service consumer and business driven - Services should be delivered at a level of granularity and abstraction that is meaningful to the Service Consumer. Driving service identification and service modeling from business and its requirements (rather than IT), will lead to the right level of granularity and abstraction. No one size fits all. The SOA and Analytics governance will improve granularity level over the life cycle of Analytical Services.

Desirable, Holistic and Sustainable - The desirable service is directly related to business oriented services and it has right level of granularity. Holistic nature comes when the service is complete and sustainable nature comes when the services meets consumer demands for security, performance and business activity monitoring.

Hide the operational & data complexity and business logic - Hiding complexity yields a gain in manageability, as well as flexibility and reuse. It is critical to model the service and address complexity as part modelling. Use the rules or policy based implementation to hide the business logic from the service consumer.

Define the service boundary – It is critical to define the service boundary both technical and business context such as application, system, data, business unit, purpose and interacting personas – customer, partners, employees. When define the service boundary you would have opportunity to engage with right technology including SAS Technology for your Service implementation.

SOA and Analytics Governance Framework - A governance program helps people do things the right way. A governance program is implemented using policies, process and organization. In the case of building Analytical Services, SAS Model Manager plays huge role in the overall Governance framework. The SOA and Analytics governance framework should manage policies, facilitate processes, collect metrics and more importantly encourage good behavior and discourage bad behavior. Having a strong Governance framework is a critical success factor for SOA and Analytics.

SOA and Analytics Quality Management – To remain competitive, an enterprise should able to deploy and provision new analytical services consumers without disrupting existing business processes. On-going quality assessment and management verifies and validates the best practices of the analytical services with respect to its functional and non-functional requirements.

CONCLUSION

Service Oriented Architecture's direct business benefits are reduced costs, reduced redundancy, better consistency, and improved efficiency through its reusability and flexibility. Similarly, Analytical Services enable fact based decisions that eventually improve the business performance. When both SOA and Analytical services are combined and integrated with the operational environment, the business can reap huge benefits in flexibility and agility in achieving analytical insight. . With SAS, you can easily align the Analytics Life Cycle and SOA life cycle to build analytical services that can be accessed via both synchronous and asynchronous service end points.

REFERENCES

Best Practices for Managing Predictive Models in a Production Environment – <http://www2.sas.com/proceedings/forum2007/076-2007.pdf>

Message Queue–Based Scoring Interface – <http://www2.sas.com/proceedings/sugi30/081-30.pdf>

SAS for Real-Time Applications – <http://www2.sas.com/proceedings/forum2008/389-2008.pdf>

Creating a SAS® Model Factory Using In-Database Analytics – <http://support.sas.com/resources/papers/proceedings11/147-2011.pdf>

Service-Oriented Architectures – Going from Buzz to Business – <http://www2.sas.com/proceedings/sugi31/001-31.pdf>

ACKNOWLEDGMENTS

Appreciation is extended to the entire SAS R&D and Product Management Team. We also appreciate Rob Stephens and Donna Bennett for reviewing the paper.

RECOMMENDED READING *BASE SAS® PROCEDURES GUIDE*

- SAS Integration Technologies - <http://support.sas.com/documentation/onlinedoc/inttech/index.html>
- SAS Model Manager - <http://support.sas.com/software/products/modelmgr/index.html>
- SAS Enterprise Miner - <http://support.sas.com/documentation/onlinedoc/miner/>

CONTACT INFORMATION

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

Saravana R Chandran, Principal Product Manager
R2275, SAS Campus Drive,
Cary, NC, 27513
+1 (919) 531-2698
saravana.chandran@sas.com

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are trademarks of their respective companies.