



SAS® High-Performance Analytics

What Could You Do with Faster, Better Answers?

Transform Your Organization and Gain Competitive Advantage

WHITE PAPER

Table of Contents

Executive Summary	1
What Could You Do with Better Insights	
Derived in a Drastically Reduced Time Frame?	2
What Is SAS® High-Performance Analytics?	3
How It Works	3
What Is the Value to Your Organization?	4
How and When High-Performance Analytics	
Can Be Useful for You	5
Use-Case Scenarios: The Vision for SAS® High-Performance	
Analytics Across Industries	6
Banking and Financial Management	6
<i>Customer Relationship Dynamic Pricing</i>	6
Communications	6
<i>Addressing Near-Real-Time Shifts in Network Demands</i>	6
Government	7
<i>Finance and Tax Offices</i>	7
<i>Criminal Justice and Public Safety</i>	7
Health Care	7
<i>Predicting and Managing Health Outcomes</i>	7
Hotels and Casinos	7
<i>Near-Real-Time Revenue Optimization</i>	7
Insurance	8
<i>Ratemaking</i>	8
Retail and Manufacturing	8
<i>Finished Inventory Management</i>	8
<i>Oil and Gas Production</i>	8
Conclusion: Why SAS?	9
About SAS	10

Executive Summary

Organizations are constantly seeking more effective ways to make decisions, increasingly relying on facts derived from huge volumes of data – and the amount of data we have to analyze is expanding exponentially. Social media, sentiment data, Web logs, sensor data, transactional data, third-party data, other big data sources ... the list goes on and on. At the same time, the pressure to make better, fact-based decisions faster has never been greater.

How is it possible to meet all of the new demands when there is the need to make hundreds of decisions each day with different degrees of complexity and frequency? Successful organizations can't wait days or weeks to look at different what-if scenarios before making a decision. Decisions need to be made in minutes or hours not days or weeks.

SAS® High-Performance Analytics addresses all of these business challenges and puts you in the driver's seat to success. In a nutshell, high-performance analytics take complex jobs that routinely take hours or days and distribute the work over highly-scalable and cost-effective processors to reduce the time it takes to analyze billions of records of data, with hundreds of input variables. The analyses are delivered in a fraction of the time required by traditional computing environments. We're talking seconds and minutes, not hours and days.

Faster results and high-impact outcomes enable you to think of things that were never before possible – and instant results allow organizations to take advantage of market changes well before their competitors do. This paper introduces SAS High-Performance Analytics, and explains how this new offering from SAS and our partners EMC Greenplum and Teradata can help you:

- Shrink the time from model inception to model deployment so you can respond more quickly to changing market conditions. Your organization will be able to differentiate and be more innovative.
- Incorporate more, new and different variables in the modeling process faster. You can exploit big data to capture value and gain competitive advantage.
- Quickly test multiple scenarios and receive accurate recommendations for transformational, game-changing decisions.

What could you do with faster, better answers? Read more to find out how SAS High-Performance Analytics works (with a database appliance) and the value it can bring to your organization. Use case scenarios illustrate how organizations across all industries can use very fast game-changing analytics to stay ahead of their markets and achieve competitive advantage.

Faster results and high-impact outcomes enable you to think of things that were never before possible – and instant results allow organizations to take advantage of market changes well before their competitors do.

What Could You Do with Better Insights Derived in a Drastically Reduced Time Frame?

What if you are a retailer? Normally it takes you 30 hours to set prices. What would it mean to your business if you could discover optimal price points in only two hours? SAS High-Performance Analytics enables you to run millions of alternate pricing scenarios quickly so individual stores can clear inventory, improve profit margins and leave the competition wondering what happened.

What if you work at a large investment banking firm that needs to recalculate entire risk portfolios at very high speeds with hundreds of computations in a very short amount of time? What would it mean if you could reduce that time from 18 hours to 12 minutes? You could quickly determine exposure, portfolio value at risk and liquidity coverage to make fast decisions on which products should be expedited to market or removed from the market. The speed and agility could give you a competitive advantage because you could fine-tune responses to changing interest rates, exchange rates and counterparty risks.

According to Oliver Schabenberger, the lead architect for high-performance analytics at SAS, these are just some of the types of problems organizations are seeking to solve more quickly. "Our customers have approached us with analytical jobs that are very long running. There are jobs that take 30 hours, 10 hours, five hours, and we need them to run faster. But we didn't just want to make something go twice as fast – we wanted to make performance gains that were transformational. This would enable our customers to actually change the way they do business. From days to hours; from hours to seconds -- this can be game-changing for the way organizations operate every single day."

If it takes a financial services company two or three days to evaluate the risk in a portfolio, they are already behind the curve. How would your business change if that risk evaluation could be done on an intraday basis, exactly when you need it?

Or, for the retailer who has to wait an entire weekend to recalculate markdown prices and tailor prices to individual stores and shopper profiles, there is a much higher chance that the operation will fail to improve profit margins or clear the way for fresh merchandise when needed. If that time could be reduced to just a few hours, there is much less operational risk.

All analysts know that creating analytical models and running the models against large volumes of data is an iterative process that takes time. If the process can be sped up, it increases the number of iterations that can be run and that increases accuracy. High-performance capabilities provide analysts the opportunity to go back and re-examine their models. They can ask, "Could we have produced better results? Could these models be improved? What if I moved different variables in or out?" When every job takes so long to run, there is no opportunity to do these things. By bringing the analytic processing time down, by making a job that runs in 10 hours run in 50 seconds, analysts have more time to work with the models to improve them.

"... We wanted to make performance gains that were transformational. This would enable our customers to actually change the way they do business."

*Oliver Schabenberger,
Lead Architect for High-Performance
Analytics at SAS*

All analysts know that creating analytical models and running the models against large volumes of data is an iterative process that takes time. If the process can be sped up, it increases the number of iterations that can be run and that increases accuracy.

SAS High-Performance Analytics also enables you to solve complex problems and derive analytic insights in minutes, not hours. This enables you to prepare, explore and model multiple scenarios using complete data, instead of only small subsets of data, and this too improves accuracy, enabling better decision making.

Furthermore, reducing the processing time needed to create vital information frees up time to ask and answer new questions. New questions drive new ideas. New ideas drive true innovations, which ultimately drive game-changing results. Why settle for good-enough decisions when you can now make better decisions even faster?

What Is SAS® High-Performance Analytics?

SAS High-Performance Analytics includes a set of analytical capabilities that can be executed in a highly scalable, in-memory distributed architecture. It allows customers to prepare, explore and model multiple scenarios using data volumes never before possible, and it provides much faster processing for complex analytical algorithms – both of which deliver better answers faster to those who need them for decision making.

The combination of high-performance SAS software and a highly tuned database appliance designed specifically for fast analytical computations serves as a scalable and reliable analytics infrastructure, enabling organizations to:

- Use multiple, complex analytical algorithms against large volumes of data to solve difficult problems.
- Derive insights at breakthrough speeds for high-value and time-sensitive decisions.
- React more quickly and confidently to seize new opportunities or manage risk in order to make better choices ahead of competitors.

How It Works

SAS High-Performance Analytics is appliance-enabled software that can be applied by organizations to help process big data and complex algorithm computations faster. The software is installed on specifically configured hardware from our database partners (EMC Greenplum and Teradata), which enables complex analytic processing to run in a distributed computing environment using in-memory processing resources.

The database software and SAS High-Performance Analytics software are loaded onto the database appliance. For optimal performance, data is pulled and collocated within the memory of the dedicated appliance for analytic processing. Because the data is stored locally in the database, it can be pulled quickly into memory again for future analysis in a very rapid manner. This enables massively parallel computing of SAS Analytics. SAS products currently enabled for high performance are Base SAS, SAS/STAT® (for high-performance statistical analysis), SAS/ETS® (for high-performance econometrics and time-series analysis) and SAS® Enterprise Miner™ (for high-performance data mining).

Why settle for good-enough decisions when you can now make better decisions even faster?

“SAS High-Performance Analytics can turn any data, including big data assets, into quicker, better decisions and ultimately competitive advantage.”

Dan Vesset, Program Vice President of IDC's Business Analytics Research
Oct. 26, 2011 SAS press release,
SAS High-Performance Analytics:
Big data answers in minutes, not hours

What Is the Value to Your Organization?

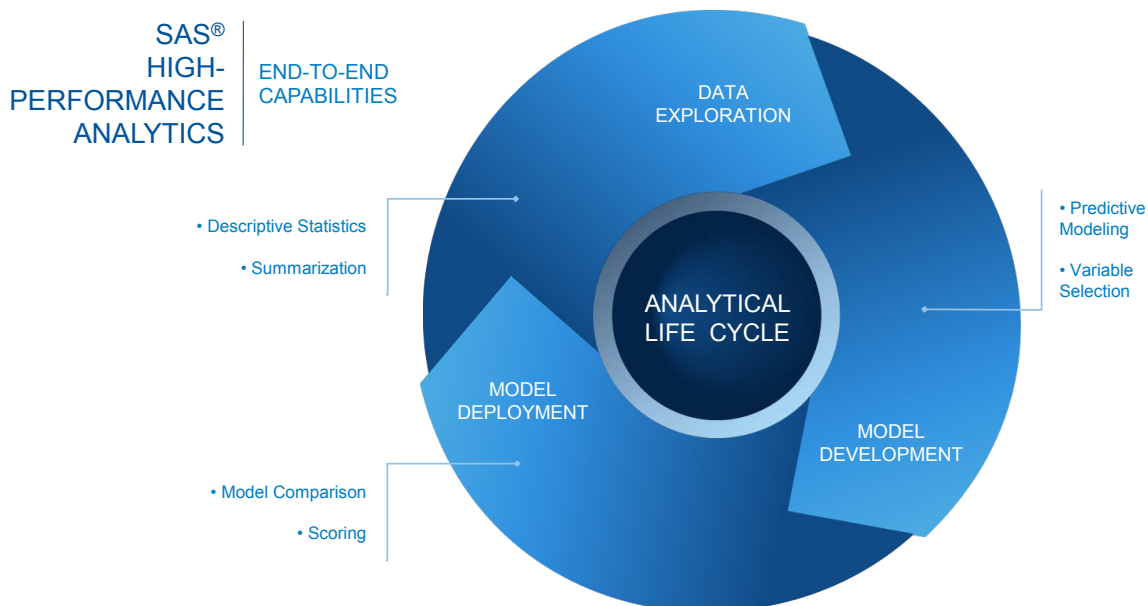
At SAS, we have always used sophisticated analytics to solve our customers' most difficult problems. As the data grows, we don't want to change the kind of analytics our customers perform. We don't want to "dumb things down" just because there is more data to crunch.

SAS CEO Jim Goodnight says, "In fact, we want to go the other way. If our customers have more data, maybe they should be asking more interesting questions. The fact that with SAS High-Performance Analytics we have been able to perform complicated analytical tasks in significantly less time provides great value to the business. Where you once ran a couple of hundred observations, you can now just as easily run the analytical processes with hundreds of billions of observations. The analysis you once did with five or 10 variables, you can now do with hundreds of variables. This is what we want to enable our customers to do. We want them to find that better life – to find the hidden value easier and faster by performing analytics faster. We want organizations to reset how they think about solving problems."

According to Goodnight and Schabenberger, the key is bringing the application as close to the data as you can. This speeds up the iterative business process of data acquisition, data analysis, variable selection, modeling and model assessment. SAS High-Performance Analytics can reduce this iterative process from four weeks to one week, or from 18 hours down to 15 minutes. If you can reduce a computation from 10 hours to 50 seconds, you have the ability to get a lot more what-if questions answered almost instantly. Instead of waiting long periods of time for one answer, you can test many models in a short time and find the right answer quickly. Goodnight says, "With SAS High-Performance Analytics, you can perform more analysis on your data in a few minutes than you could have done previously in a week."

"We want organizations to reset how they think about solving problems."

Jim Goodnight,
SAS CEO



With SAS® High-Performance Analytics, you can greatly speed up the entire iterative analytical life cycle. This enables analysts to test more theories, ask more what-if questions and quickly produce more accurate results.

How and When High-Performance Analytics Can Be Useful for You

- Do you find yourself forced to limit the amount of data you analyze? Are you not able to take full advantage of *all* your data?
- Are you limited to asking simple questions? Or not asking questions at all?
- Are you getting partial answers too late to make a difference?
- Are your analysts forced to compromise on their use of modeling techniques? For example, do they have to perform a simpler linear regression instead of a better-fitting logistic regression because the linear regression processing fits within an IT-required time frame?
- Do your analysts have to forego model testing and refinement because there is not enough time to perform multiple iterations?

If you answered yes to any of the above questions, your organization could gain value from high-performance analytics. With SAS High-Performance Analytics, organizations can capitalize on opportunities for growth that they could not otherwise find, improve performance on an ongoing basis and differentiate themselves in the marketplace. It enables decision makers to ask difficult questions more frequently and make market entry and exit decisions quickly and confidently. Armed with this information, organizations can gain an edge on the competition, and it can mean the difference between being a market leader and a follower.

Finer and more accurate results enable organizations to realize significant business value, and drive revenue and bottom-line cost savings. The blazingly fast performance of SAS High-Performance Analytics can help organizations evaluate alternate scenarios and detect changes in volatile markets in time to make the best recommendations. In minutes, not hours, a bank can quickly and accurately determine the level of credit exposure in its consumer-lending portfolio and reduce charge-offs. Too much credit exposure can lead to high default rates and charge-off percentages. One financial institution has been able to calculate the probability of a loan default in 84 seconds versus 167 hours! Even a 1 percent reduction in loan charge-offs can result in huge savings for a global bank.

SAS High-Performance Analytics can help organizations fully exploit all data, instead of just aggregates or subsets of data, thereby improving accuracy for more targeted, high-impact decisions. Users can handle more variables, employ more complex modeling techniques, perform model iterations more frequently and test new ideas. Users will spend more time doing important analyses instead of developing a statistically significant sample to produce possibly similar results. For example, thousands of segment-specific models can be created to incorporate unique customer preferences and behaviors, instead of just applying a general model to all data. This can identify top predictors of those likely to respond to a campaign across millions of highly qualified prospects.

What Are the Differences Between High-Performance Analytics and Other High-Performance Computing Options?

Grid computing: Enables SAS to automatically use a centrally managed pool of resources assigned in a grid environment to achieve workload balancing, high availability and parallel processing.

In-database analytics: In-database computing moves relevant tasks closer to data, running computations inside the database to promote better data governance and faster results. SAS Analytics run inside the database, thereby avoiding time-consuming data movement and conversion. However, analytic processing is still constrained by what the database workload management system allocates to it.

In-memory analytics: Divides up analytic processes into easily manageable pieces. In this case, the computations are distributed in parallel across a dedicated set of blade servers.

High-performance analytics: A relatively new model that uses in-memory analytics processing. Instead of running the analytics as part of the database process, the analytics run on the same hardware as the database, but as a separate process, quickly passing data back and forth as needed, allowing for much more complex analytics than in-database. SAS Analytics processing is not constrained by the database workload management system and therefore can make use of all system resources.

Use-Case Scenarios: The Vision for SAS® High-Performance Analytics Across Industries

SAS High-Performance Analytics can provide true value across all industries. Here are just a few possibilities. Imagine what your organization could do with better, faster answers.

Banking and Financial Management

Customer Relationship Dynamic Pricing

Many banks today are hoping to grow consumer and small business revenues through cross-sell and up-sell techniques that increase the number of products per customer and also create “sticky” relationships that reduce attrition. The challenge in selling additional products to each customer is there may be no way to modify the price of new products under consideration based on the current value of the relationship and how this could change with the addition of other new products. With high-performance analytics, the bank representative could assess the customer’s current use of existing bank products and services along with associated profitability and combine that information with in-house propensity, credit scores and external data (such as outstanding loans and other financial relationships) to gain a view of the customer’s potential lifetime value. The overall value to the bank through the addition of high-performance analytics is that every customer interaction can be based on optimizing the price of new products for each customer in a way that increases retention, grows revenue and improves the bank’s profits while providing the optimal customer experience for each individual consumer or business client.

The overall value to the bank through the addition of high-performance analytics is that each customer interaction can be based on optimizing the price of each new product for each customer in a way that grows revenue and optimizes the bank’s profits while providing the optimal customer service for each individual consumer or business.

Communications

Addressing Near-Real-Time Shifts in Network Demands

Network operators no longer have the luxury of taking months to address changes in network demand. Bandwidth consumption is becoming harder to forecast as smartphones, tablets, on-demand HD video and cloud computing can create instant shifts in consumption patterns. Traditional capacity planning was developed for an era when it was acceptable to spend months upgrading network elements to meet increased demand. Rising consumer expectations require new methodologies that correlate customer data with near-real-time network status information. A large communications network executes and tracks billions of events each day involving millions of customers. High-performance analytics enables nearly instantaneous analysis of network performance and the ability to measure the effect on individual customers. Correlating customer data with near-real-time network performance monitoring enables the network operator to know not just which parts of the network have degraded service, but which customers are affected and what corrective actions make the most business sense. By dynamically reacting to customer demands, network performance operators can deliver optimal performance all of the time.

High-performance analytics can analyze a large complex network in near-real time, enabling the network to dynamically reconfigure itself to provide the optimum performance at all times.

Government

Finance and Tax Offices

Tax offices are often asked to answer questions for government officials (i.e., If you increase income taxes by one point, what is the impact on the budget?). Many revenue statistics departments use SAS to model and assess the implications of proposed budget options and present their findings. SAS is also used to do simulations to find relations between the macro economy and tax collection. SAS can be used to improve a city or country's financial status by identifying and prioritizing uncollected debt that is owed, as well as identifying new fraudulent tax filings. The more information (both structured and unstructured text data) that can be analyzed quickly allows tax organizations to identify fraudulent cases sooner; simulate the effects on a tax code faster; run scenarios to quantify the impact of changing the macroeconomic factors on the health of the community; and capture information regarding the likelihood of payment through various models when changing tax codes.

The more information (both structured and unstructured text data) that can be analyzed quickly allows tax organizations to identify fraudulent cases sooner ...

Criminal Justice and Public Safety

Law enforcement organizations are not only responsible for detecting, preventing, responding to and solving crimes against people and property in their communities, they are also tasked with a more global approach to policing that challenges resources, skills and personnel limits. With high-performance analytics, law enforcement personnel can more quickly identify and focus on crime hot spots, enabling police to respond more rapidly and counter criminal activity more readily. They are also able to integrate vast amounts of data, both structured and unstructured (social media, text, etc.), to anticipate and prevent criminal activity. Officer and community safety can be enhanced if investigators can move from a reactive to a more proactive approach.

Health Care

Predicting and Managing Health Outcomes

Personalized medicine promises the ability to optimize health outcomes based on genetic information for each patient. Of course, genetic information is massively data intensive (a single sequencer can generate terabytes of data per week). High-performance analytics offer the possibility to analyze what treatment options are best for a particular patient based on their personal genetic makeup, rather than just comparisons to other patients similar to them. The ability to analyze big data extremely quickly means high-performance analytics could move personalized medicine research from the lab to the clinic so it could be used to drive individual treatments.

The ability to analyze big data extremely quickly means high-performance analytics could move personalized medicine research from the lab to the clinic so it could be used to drive individual treatments.

Hotels and Casinos

Near-Real-Time Revenue Optimization

An average-size hotel of 250 rooms needs to calculate pricing and inventory controls for at least 365 days into the future for approximately five room types and eight lengths of stay. This results in 3.65 million data points to review on a daily basis. With high-

performance analytics, rather than produce pricing and inventory control decisions using a batch process, it is possible to optimize them in near-real time, even after each reservation is made. Increased benefits include maximized revenue for every decision (instead of a group of decisions) and the ability to respond much faster to changing trends.

Insurance

Ratemaking

Product pricing or ratemaking is the process of establishing rates charged by an insurer for accepting the risk. Insurance companies, and specifically actuaries, rely heavily on using historical data to predict future behavior or create premium rates to price products. In the past, actuaries have relied on univariate or one-way analysis for pricing and monitoring price efficiency. However, today more insurers are using advanced analytical techniques such as generalized linear modeling. Also in the past, actuaries have often relied on using a subset of historical data to run pricing models because the time it takes to prepare the data and run the models has been too time-consuming. To combat these problems, insurers are now turning to high-performance analytics to provide faster processing on the growing volumes of available data.

Retail and Manufacturing

Finished Inventory Management

Retailers, manufacturers or any organization with significant volumes of inventory in a large variety of locations and distribution levels can benefit from SAS High-Performance Analytics. Some inventory allocation decisions are time critical, and significant revenue loss and decreased customer satisfaction can occur by not having the right item in the right place at the right time. While most demand forecast prediction is done on historical shipping data, the availability of point-of-sale (POS) data makes it possible to perform in-memory processing of historical plus time-latency-adjusted POS data on very large data sets.

Oil and Gas Production

SAS High-Performance Analytics can bring increased value to oil and gas production by scoring alert data that streams in from platforms and wellheads. The data alerts come from hundreds of sensors at rates of up to every millisecond. The flow of data is heavy, constant and intense. SAS High-Performance Analytics produces near-real-time scoring and surfaces the highest score for immediate action, potentially saving hundreds of millions of dollars and preventing catastrophic failures and damaged brand reputation.

... Rather than produce pricing and inventory control decisions using a batch process, it is possible to optimize them in near-real time, even after each reservation is made.

... Insurers are now turning to high-performance analytics to provide faster processing on the growing volumes of available data.

SAS High-Performance Analytics produces near-real-time scoring and surfaces the highest score for immediate action, potentially saving hundreds of millions of dollars and preventing catastrophic failures and damaged brand reputation.

Conclusion: Why SAS?

- **Get integrated technology from the leaders in business analytics and database appliances.** SAS' expertise in business analytics and our partners' (EMC Greenplum and Teradata) deep understanding of massively parallel processing architecture offers our customers an integrated software and hardware appliance designed to tackle large, difficult problems at a much faster pace.
- **Not all in-memory approaches are created equal.** SAS provides the only in-memory analytics offering that delivers high-end predictive analytics and uses big data to produce time-sensitive insights. In-memory analytics are a hot topic right now, but not all in-memory approaches are created equal. SAS High-Performance Analytics is not just about using query, reporting and descriptive statistics within an in-memory environment. It is truly about applying sophisticated, world-class analytical, forecasting and optimization techniques to quickly solve complex, real-world business problems.
- **Address the entire model development and deployment life cycle.** Unlike other competitors, SAS High-Performance Analytics can perform analyses that range from descriptive statistics and data summarizations to model building and scoring of new data at breakthrough speeds. The fast results help our customers run more models, ask more questions, extract more value from the data and stay ahead of the competition.
- **Discover, analyze and process big data that cannot be dealt with using traditional approaches.** The adoption of this new class of technology extends SAS' leadership in the analytics market – allowing you to tackle big data, analyze more complex scenarios and answer real-world questions. Shrink the time it takes to go from model inception to deployment. SAS High-Performance Analytics delivers blazing-fast performance so you can evaluate numerous scenarios and quickly detect, and act on, changing market conditions and achieve competitive advantage.

Still wondering if SAS High-Performance Analytics can provide real value for your organization? Ask the industry analysts who have routinely listed SAS as a leader in business analytics. SAS has long been supporting big data environments and we specialize in solving the complex challenges our customers face – the kind that can transform your business. Join our more than 55,000 customer sites and tackle your big data challenges, solve your most difficult business problems and increase the speed and quality of your fact-based decision making with SAS High-Performance Analytics.

“Harnessing the power of all those data points in parallel lets you do things you never thought about before because the data was just way too big. We’re having a lot of fun right now working on these kinds of problems.”

Jim Goodnight,
SAS CEO

Learn more

about SAS software and services for high-performance analytics:
sas.com/hpanalytics.

About SAS

SAS is the leader in business analytics software and services, and the largest independent vendor in the business intelligence market. Through innovative solutions, SAS helps customers at more than 55,000 sites improve performance and deliver value by making better decisions faster. Since 1976, SAS has been giving customers around the world THE POWER TO KNOW®. For more information on SAS® Business Analytics software and services, visit sas.com.



SAS Institute Inc. World Headquarters +1 919 677 8000

To contact your local SAS office, please visit: sas.com/offices

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. Copyright © 2012, SAS Institute Inc. All rights reserved. 105592_S88388_0212