



---

## Highlights

IBM Netezza and SAS together provide appliances and analytic software solutions that help organizations improve operational efficiency, strengthen customer relationships, increase profitability and reduce risk.

- **Integrated technology** – SAS solutions deployed on IBM Netezza data warehouse appliances produce analytic results dramatically faster than previously attainable.
- **Business value** – Organizations can take advantage of query performance to conduct more analysis on larger data sets while enabling real-time analytics to deliver new business processes.
- **Benefits** – Improved efficiency and increased productivity derived from high-performance analytics lead to competitive advantage.

# High-Performance Business Analytics: SAS and IBM Netezza Data Warehouse Appliances

Today's organizations are faced with diverse issues, a wider range of regulations and heightened global competition. Data volumes are growing exponentially, challenging an organization's ability to capture, store and analyze business-critical information. There has never been a greater need for making decisions based on empirical data and being able to adjust strategies quickly. By delivering insights that are garnered from data about customers, suppliers, operations, performance and more, organizations can solve complex problems, manage performance to achieve measurable objectives, achieve sustainable growth through innovation, and anticipate and manage change.

It is not easy to aggregate, analyze and process large volumes of data and analytic models quickly enough when business decisions must be made in hours, minutes or even seconds. IBM, partnering with SAS®, offers the IBM® Netezza® data warehouse appliance family, optimized to run SAS solutions.

## Real business impact, right from the start

For more than 30 years, SAS — a world leader in business analytics software — has worked with IBM to deliver solutions that help organizations improve operational efficiency, strengthen customer relationships, increase profitability and reduce risk. Across a full-range of industries, organizations have implemented SAS and IBM technologies to capitalize on proven solutions that can transform their business.

IBM, with the IBM Netezza data warehouse appliance, and SAS partner to deliver a platform optimized for high-performance business analytics. Together IBM and SAS are putting their deep knowledge and extensive experience to work, with the development of tightly integrated business analytics offerings.



---

## Solutions from SAS that directly leverage the IBM Netezza data warehouse appliance include:

- **SAS/ACCESS® Interface to IBM Netezza data warehouse appliance** — an out-of-the-box solution that provides direct connectivity between SAS and the IBM Netezza data warehouse appliance.
  - **SAS Enterprise Miner™** — streamlines the data mining process to create highly accurate predictive and descriptive models based on large volumes of data from across the enterprise.
  - **SAS Scoring Accelerator** — enables customers to publish scoring models created in *SAS Enterprise Miner* into database-specific functions to be deployed and then executed directly within the IBM Netezza data warehouse appliance.
- 

## High-performance data warehousing and analytics

IBM Netezza data warehouse appliances are purpose-built to make advanced analytics simpler, faster and more accessible. They deliver the proven performance, value and simplicity organizations need to leverage their growing data and make fact-based decisions. IBM Netezza data warehouse appliances offer a powerful, cost-effective, easy-to-install way for organizations to scale their business intelligence and analytical infrastructure.

A key performance advantage of the IBM Netezza data warehouse appliance family comes from its patented Asymmetric Massively Parallel Processing (AMPP™) architecture. The AMPP architecture is a blade-based streaming architecture that uses commodity blades and storage. It is combined with IBM's patented data filtering using Field Programmable Gate Arrays (FPGAs) to deliver blisteringly fast query performance and modular scalability on highly complex mixed workloads. It supports tens of thousands of BI, advanced analytics and data warehouse users.

IBM Netezza data warehouse appliances are designed to minimize time-consuming, ongoing administrative tasks needed to maintain traditional databases. Eliminated tasks include query indexing, storage management, buffer pool tuning, memory allocation and schema optimization. This results in a reduction in the dependence on the IT staff, a cut in total cost of ownership and a shorter time-to-value.

## Integrated solutions for in-database processing

Business analytics solutions are only valuable if they can deliver the right information, to the right person, at the right time. *SAS/ACCESS Interface to Netezza* offers optimized read/write connectivity and allows a fast connection for loading and extracting data to and from SAS tools, platforms and solutions. It allows SAS to read/write directly to IBM Netezza tables, and provides implicit SQL support which generates IBM Netezza specific SQL.

Now analysts and modelers can take all the intelligence of SAS scoring algorithms and embed them into an IBM Netezza data warehouse appliance.

*SAS Scoring Accelerator* for IBM Netezza data warehouse appliances deploys models quickly, reduces data movement and leverages the power of the IBM Netezza data warehouse appliance. It combines the statistical models quickly, reduces data movement and leverages the power of the IBM Netezza data warehouse appliance. *SAS Scoring Accelerator* combines the statistical transformation and modeling methods available in *SAS Enterprise Miner* with the scalability and processing speed offered through the database system. It takes the complete scoring model code, the associated property file that contains model inputs and outputs, and a catalog of user-defined formats, and deploys (or publishes) them within the data warehouse database. This results in a reduction of overall cost because the scoring function is moved to the database, and the common security, auditing and administration capabilities offered by the database are honored and leveraged.

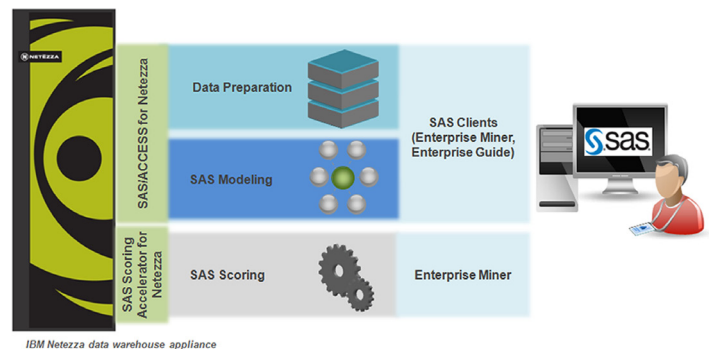
*“Every SAS client who adds Netezza drastically reduces their execution times. 20 hour jobs reduced to 20 minutes is the standard.”*

— Dana Rafiee, U.S. Managing Director, Destiny Corp.

*SAS Scoring Accelerator* for IBM Netezza data warehouse appliances allows analysts to publish scoring models created in SAS Enterprise Miner into database-specific functions to be deployed and then executed directly within the database environment. In-database scoring helps to generate results quickly, reduces data movement and leverages the power of the IBM Netezza data warehouse appliance. *SAS Scoring Accelerator* automates the scoring of models created in SAS Enterprise Miner with the scalability and processing speed offered through the database system. It takes the complete scoring model code, the associated property file that contains model inputs and outputs, and a catalog of user-defined formats, and deploys (or publishes) them within the data warehouse database. This results in a reduction of overall cost and efficient utilization of resources because the scoring function is moved to the database, and the common security, auditing and administration capabilities offered by the database are honored and leveraged.

- Optimizes SAS Scoring
- Scales to manage mixed workloads
- Achieves faster model improvement cycles
- Reduces data movement, latency and complexity of the scoring process
- Achieve higher model-scoring performance and faster time to results
- Increases analysts' productivity
- Consolidates data used for scoring to improve regulatory compliance

As an alternative to an extract/score/reload operation, users can write a program from scratch in C, C++ or another supported language and run the scoring job for use on an IBM Netezza data warehouse appliance. Custom-coding a scoring model from scratch takes time to test, validate and tune. Customers report cycle times of three to six months to deploy a scoring model into production. Manual coding also introduces a source of error into the process, so that scoring jobs must be exhaustively validated to ensure they produce the same results as the original predictive model. *SAS Scoring Accelerator*, deployed in IBM's massively parallel architecture, significantly reduces the time it would take to run in a single-threaded out-of-database operation for a custom scoring job. Thus, it provides the benefit of a scalable and high-performance scoring function, without the costs and delay.



## Business value

**Accelerated analytics time-to-results:** SAS technology offers a solution that makes it possible for organizations to complete model development in less time, while realizing higher model-scoring performance. This allows for the deployment of analytic models more quickly by reducing unnecessary data movement and making the most of database scalability. This yields more-timely information for business decision making. It also significantly reduces data movement and latency, improving the accuracy and effectiveness of analytic models.

---

## Release notes

As of Release 9.3, SAS supports seven in-database procedures (PROCs) for use on IBM Netezza data warehouse appliances: **FREQ**, **MEANS**, **RANK**, **REPORT**, **SORT**, **SUMMARY** and **TABULATE**. These PROCs are commonly used for basic management reports (such as daily response reports, sales reports, staffing reports, survey crosstabs, and so forth). Since SAS datasets can be held in the IBM Netezza data warehouse appliance, SAS user groups can use these in-database PROCs to set up a high-performance reporting environment that runs entirely in-database on an IBM Netezza data warehouse appliance.

---

**Strengthened data governance:** With *SAS Scoring Accelerator* for IBM Netezza data warehouse appliances, modelers determine where the data is — and what form it is in — at any time. This allows the client to use database security for data privacy and to achieve improved compliance with corporate governance requirements.

## About SAS

SAS is a leader in business analytics software and services, and the largest independent vendor in the business-intelligence market. Through innovative solutions, SAS helps customers improve performance and deliver value by making better decisions faster. SAS users receive a full suite of support services at no extra charge, including skilled telephone technical support and unlimited, around-the-clock online technical support. Visit [www.sas.com](http://www.sas.com) to learn more.

## About IBM Netezza data warehouse appliances

The IBM Netezza data warehouse appliance revolutionized data warehousing and advanced analytics by integrating database, server and storage into a single, easy to manage appliance that requires minimal set-up and ongoing administration while producing faster and more consistent analytic performance. The IBM Netezza data warehouse appliance family simplifies business analytics dramatically by consolidating all analytic activity in the appliance, right where the data resides, for blisteringly fast performance. Visit [netezza.com](http://netezza.com) to see how our family of data warehouse appliances eliminates complexity at every step and lets you drive true business value for your organization. For the latest data warehouse and advanced analytics blogs, videos and more, please visit: [thinking.netezza.com](http://thinking.netezza.com).

## IBM Data Warehousing and Analytics Solutions

IBM provides the broadest and most comprehensive portfolio of data warehousing, information management and business analytic software, hardware and solutions to help customers maximize the value of their information assets and discover new insights to make better and faster decisions and optimize their business outcomes.



---

© Copyright IBM Corporation 2011

IBM Corporation  
Software Group  
Route 100  
Somers, NY 10589  
U.S.A.

Produced in the United States of America  
October 2011

All Rights Reserved

IBM, the IBM logo, and [ibm.com](http://ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at [ibm.com/legal/copytrade.shtml](http://ibm.com/legal/copytrade.shtml)

Netezza is a registered trademark of Netezza Corporation, an IBM Company.

Other company, product and service names may be trademarks or service marks of others.

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.



Please Recycle