

Solution Guide: IBM® Power Systems™ and IBM Storage for SAS® Viya®

Empower SAS Analytics

WHY INFRASTRUCTURE MATTERS FOR SAS ANALYTICS?

Organizations leverage analytical models to make data driven, real time decisions. But many of them struggle in operationalizing these models to make decisions at scale. Analytical solutions from SAS reduce time to value by eliminating redundant steps in the modeling life cycle and supporting cohesion across the information chain from data discovery to decision management.

SAS Viya, a cloud-enabled, and open analytics engine, extends the SAS analytics platform to enable decision makers, data analysts and developers to collaborate and realize faster results. Across the analytics lifecycle – from data discovery, model development and deployment, to operationalizing analytics – SAS Viya addresses a wide range of compute demands depending on the type of applications that run on it. Some are very low on compute but high on data volume, some aggregate data from various sources, and others have AI, ML/DL capabilities and require GPUs.

SAS Viya's broad range of analytics capabilities require compute and storage infrastructure that can handle the complexity of the workload or else, clients carry the risk of reduced performance, reliability and scalability. Extracting the right insights from any analytics tool is an iterative process, often involving several rounds of slicing and dicing of data sets. The data scientist or LOB analyst depends on fast data processing to explore data and quickly generate insights. Choosing the right SAS infrastructure is therefore critical.



EMPOWER SAS ANALYTICS WITH IBM POWER SYSTEMS AND IBM STORAGE

IBM and SAS have been collaborating since the SAS's founding to provide best in class compute and storage solutions that will help organizations fully leverage the capabilities of SAS analytics. IBM Systems empowers SAS analytics by enabling clients to accelerate insights with industry leading performance. This agile, full stack solution delivers the wide range of compute needed across the analytics lifecycle and handles the massive volume and velocity of data coming into the SAS platform with maximum resilience.

1. ACCELERATE INSIGHTS

Accelerating performance of SAS analytics jobs is critical for business executives who need insights to make real-time decisions and data analysts who need to iterate to develop models that deliver the right insights. IBM Power Systems and IBM Storage are designed from the ground up for big data analytics. IBM solutions handle massive data throughput with industry leading I/O and memory bandwidth. This has been a key differentiator for SAS 9 on AIX and is now even more important with Viya's in-memory, parallel load approach that increases data throughput. The compute demands of SAS workloads vary across the phases of the analytics lifecycle. Similarly, they vary based on the type of SAS workloads. For example, compute demands for event stream processing differ from those for visualizing data. The ability to scale and deploy more threads per core allows IBM Power Systems to tackle this mix of computing demands at the scale and performance required to help clients meet their SLAs.

With an increased presence of AI in SAS's portfolio from machine learning based forecasting and optimization, to computer vision and natural language processing (NLP), the value of GPU-enabled systems in SAS environments is growing. IBM Power Systems AC922 – the GPU-enabled system designed for Enterprise AI (and [the foundation of the US Dept of Energy's Summit and Sierra supercomputers](#), the two most powerful supercomputers in the world¹) – is engineered to be the most powerful AI training for platform accelerating insights from SAS analytics.

IBM Storage also plays a critical role in faster transformation of data into insights. IBM FlashSystem 9100 is an end to end NVMe accelerated solution that provides petabytes of data storage. It delivers throughput of up to 34GB per second in 2U and IOPS of up to 3.75 million IOPS per 2U.

IBM POWER9

The processor that drives world's most powerful supercomputers.

4x²


More threads for high performance cores vs compared x86

5x+

More I/O bandwidth than compared x86

1.8x

More memory bandwidth per socket vs compared x86



NVLink 2.0
PCIe Gen4

2. AGILE, FULL STACK SOLUTION

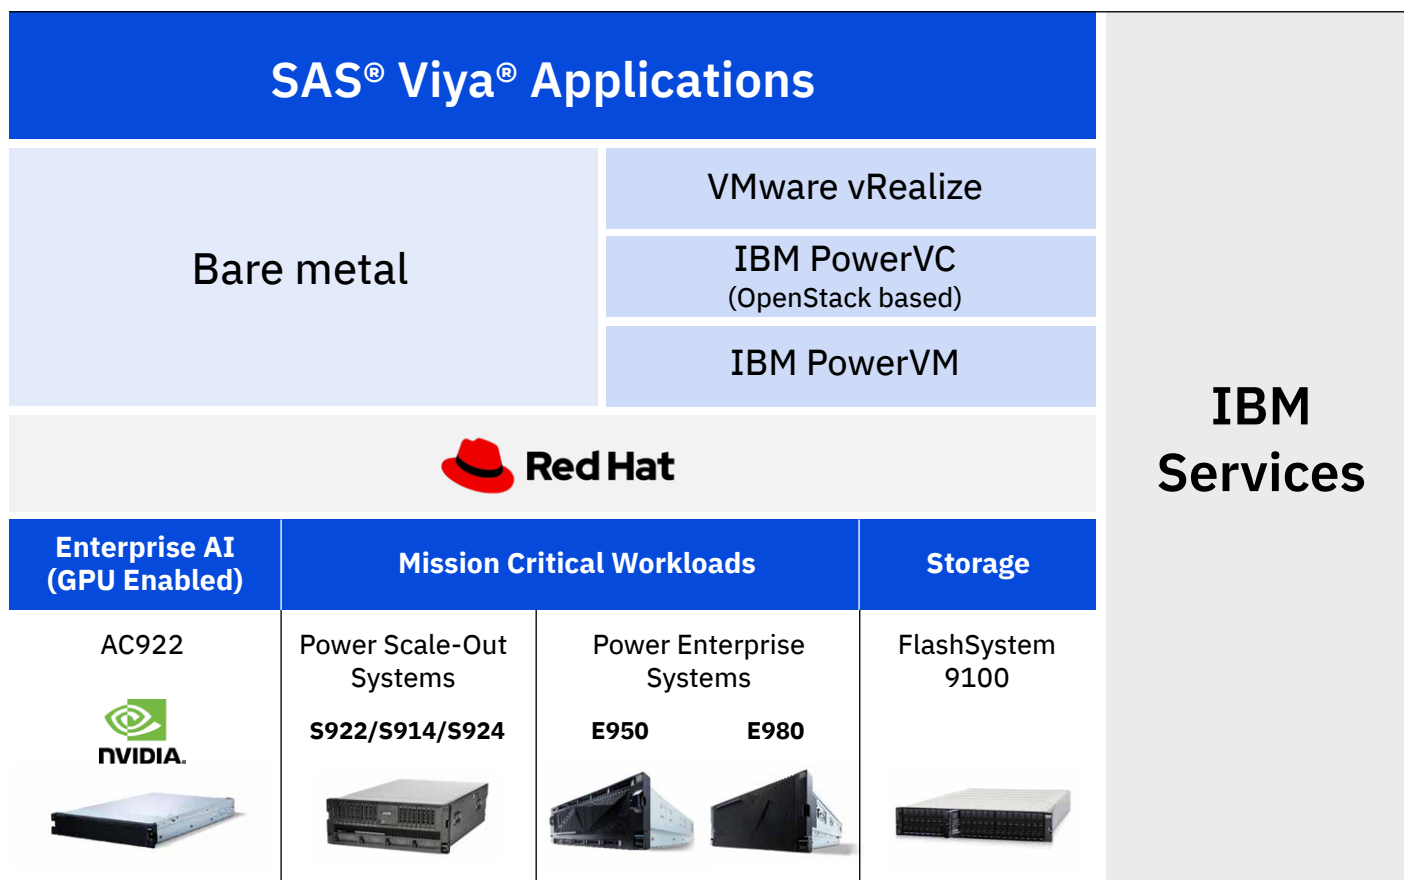
IBM helps clients build and optimize their SAS solutions with full stack infrastructure solutions. From servers to storage, to Red Hat Enterprise Linux, and services and support, IBM offers end-to-end infrastructure solutions that provide flexible deployment options to tackle a mix of compute demands, enable on-demand scaling and simplify management.

¹Top500 supercomputers, June 2019 - <https://www.top500.org/lists/2019/06/>

²4X - P9 cores provide 4x threads of x86 cores

5X+ - 5.6x I/O bandwidth claim based on NVIDIA measurement test conducted on a Xeon E5-2640 V4 +P100 vs Power9 + V100 (12 GB/s vs 68 GB/s rated)

1.8X - bandwidth is based on 230 GB/sec per socket for POWER9 and 128GB/sec per socket for x86 Scalable Platform Intel



Flexible deployment options

The right infrastructure mix for SAS depends on multiple factors. You may be deploying SAS for the first time and looking for a bare metal cluster to run SAS VDMML in Massive Parallel Processing (MPP) mode or may want to run SAS Viya in your existing POWER9 system alongside SAS 9.4 in Symmetric Multi-Processing (SMP) mode. No matter what the unique needs are - bare metal or virtualized, GPU or non-GPU, scale up or scale out clusters - IBM Power Systems offers the flexibility to meet them.

On demand scaling

In a virtualized environment, infrastructure footprint can be minimized by deploying multiple production workloads on Linux, IBM i and AIX in the same system. You can also deliver infrastructure as a service within the cloud with OpenStack based IBM PowerVC. This enables on-demand scaling of the environment by dynamically allocating capacity based on changing resource needs.

Simplified management

If datacenter operations are managed using VMwarevRealize platform, PowerVC integrates with the platform enabling single pane management. Managing large storage systems requires many hours of monitoring, analysis, decision-making and adjustment. When problems arise, troubleshooting complex storage infrastructure and implementing the most effective solutions can be problematic, to say the least. To address these challenges and reduce both manual labor and mistakes, IBM FlashSystem 9100 solutions come with Storage Insights, an enterprise-proven, AI, cloud-based system insights platform to help you better understand trends in storage capacity and performance and expedite resolution when support is required. Storage Insights monitors the health, capacity and performance of all IBM block storage and external storage under management on a single pane of glass, helping IBM customers understand

and plan storage capacity and performance. The program provides proactive best practices and uses AI-based analytics to help identify potential issues before they become problems. When support is needed, Storage Insights helps speed resolution by simplifying opening tickets, automating log uploads to IBM, and providing configuration, capacity and performance information to IBM technicians.

Manage and protect data

Software defined storage with IBM Spectrum Virtualize provides an ideal way to manage and protect huge volumes of data used for big data analytics. Its data reduction technologies increase the amount of data you can store in the same space by up to five times. It also helps protect data with software-based encryption while also eliminating storage downtime with nondisruptive movement of data from one type of storage to another.

3. MAXIMIZE RESILIENCE

Power Systems is ranked the most reliable server for the past 11 years by ITIC³. It has consistently delivered more than 99.999% uptime and has built-in intelligent memory protection to detect and fix potential faults before they lead to system failure. In commodity architecture, comparable technology is optional and affects performance when used. Live Partition Mobility in Power Systems enables zero impact planned maintenance by non-disruptive

mobility of workloads across generations of Power Systems and IBM HyperSwap also delivers the outcome for storage on FlashSystem 9100.

SAS Viya on IBM Power Systems and IBM Storage enable businesses to make decisions at scale. The full stack infrastructure solutions from IBM empower SAS analytics with industry leading performance, flexible deployment options and maximum resilience.

LEARN MORE AT: ibm.biz/SASViya

³ [ITIC 2019 Global Server Hardware, Server OS Reliability Survey](#):

- P9 cores provide 4x threads of x86 cores
- 5.6x I/O bandwidth claim based on NVIDIA measurement test conducted on a Xeon E5-2640 V4 +P100 vs Power9 + V100 (12 GB/s vs 68 GB/s rated)
- 1.8X bandwidth is based on 230 GB/sec per socket for POWER9 and 128GB/sec per socket for x86 Scalable Platform Intel product brief: <https://www.intel.com/content/dam/www/public/us/en/documents/product-briefs/xeon-scalable-platformbrief.pdf?asset=14606>

© Copyright IBM Corporation 2019.

IBM, the IBM logo, and ibm.com are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and

service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at ibm.com/legal/us/en/copytrade.shtml, and select third party trademarks that might be referenced in this document is available at ibm.com/legal/us/en/copytrade.shtml#section_4.

This document contains information pertaining to the following IBM products which are trademarks and/or registered trademarks of IBM Corporation: IBM®, POWER9™, Power Systems™