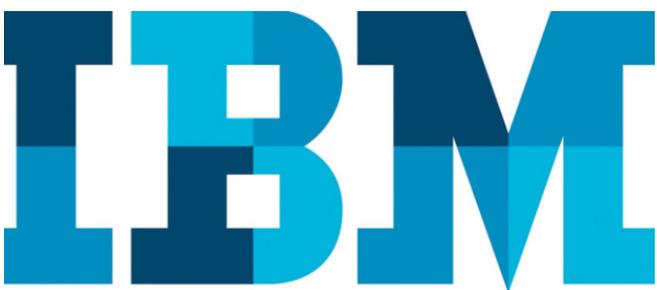


SAS Viya 3.5 deployed in an IBM POWER9 mixed environment

Table of contents

<i>Introduction</i>	2
<i>Architecture</i>	3
<i>Testing the infrastructure</i>	5
<i>Test results</i>	6
<i>Summary</i>	7
<i>Get more information</i>	7
<i>About the authors</i>	7



Introduction

Highlights

- Flexible Viya deployment across IBM Power System E980 and Power AC922 servers
 - Scalable and resilient Power E980 server
 - GPU-enabled Power AC922 for AI and ML Viya workloads for faster insights
-

The IBM® Power Systems™ portfolio of servers enables flexible deployment options for running SAS Viya. It offers flexible systems with the highest reliability¹. The enterprise and scale-out servers can run Linux® in addition to IBM AIX® and IBM i, and allow for the consolidation of multiple workloads on a single system which can reduce data center footprint. An example of this is deploying multiple SAS workloads on the same system. Customers can also take advantage of advanced virtualization features such as Capacity on Demand (CoD) to dynamically activate processor capacity and memory resources for ease of scalability and flexibility in adapting to growing demands the infrastructure.

This technical brief demonstrates the deployment of SAS Viya 3.5 operating in the massively parallel processing (MPP) mode in a mixed environment spread across a cluster of IBM PowerVM® logical partitions (LPARs) on an IBM Power® System E980 scale-up server (based on the IBM POWER9™ processor technology) and a scale-out Power AC922 server.

The goals, results, and supporting information about the solution implemented are documented in this technical brief.

Goals

- Demonstrate deployment of Viya 3.5 (in MPP mode) in a mixed environment
- Showcase the flexibility of POWER9 enterprise-class servers
- Successfully run a Viya workload for functional verification

Reference architecture

Software

- SAS Viya 3.5
- IBM PowerVM
- RHEL 7.6 LE
- Ansible 2.9

Hardware

- IBM Power E980 server
- IBM Power AC922 server

Network

- 10 GbE network
 - NVMe drives local to server
-

Architecture

The infrastructure configured was the IBM Power System E980 and AC922 server. The storage used was high-speed and low-latency front-facing NVMe drives local to the server.

A 10GbE private network was set up for internode traffic.

Software building blocks are SAS Viya 3.5 deployed on RHEL 7.6 LE for the IBM PowerPC® architecture. The test bed employed for functional verification was the SAS Viya workload.

SAS Viya and its software in a mixed MPP environment is installed across a set of LPARs on the scale-up Power E980 server and one or more scale-out Power AC922 servers. In this case, the SAS Cloud Analytics Services (CAS) controller, SAS Programming Runtime Environment and management services run in the LPARs and CAS worker nodes could run either on LPARs (for non-GPU workloads) or on scale-out servers (for CPU or GPU workloads).

Figure 1 depicts a logical representation of a mixed MPP mode Viya deployment. The Viya components are in green and non-Viya software are in the other colors. This environment provides a combined set of benefits for scaling, supporting parallel loading, and processing of data, a growth path, scaling by using unused resources (which allows adding more LPARs and resources to existing LPARs or worker nodes), simplified networking between LPARs, supporting GPU-enabled workloads as they are added, and performance benefits with NVLink 2.0 technology.

Figure 2 depicts the architecture of a typical mixed MPP mode Viya deployment in production. The storage used is internal drives local to the server [high-speed solid-state drive (SSD) or NVMe]. The Virtual I/O Server (VIOS) was used to create the shared Ethernet adapter (SEA) for all client LPARs for external access *to* and *from* the LPARs.

A slightly smaller version of the deployment in Figure 2 was set up for functional verification and the details are explained in the subsequent sections.



Example: IBM Power System E980



Example: GPU-enabled AC922 servers

Figure 1. Logical representation – MPP mixed environment

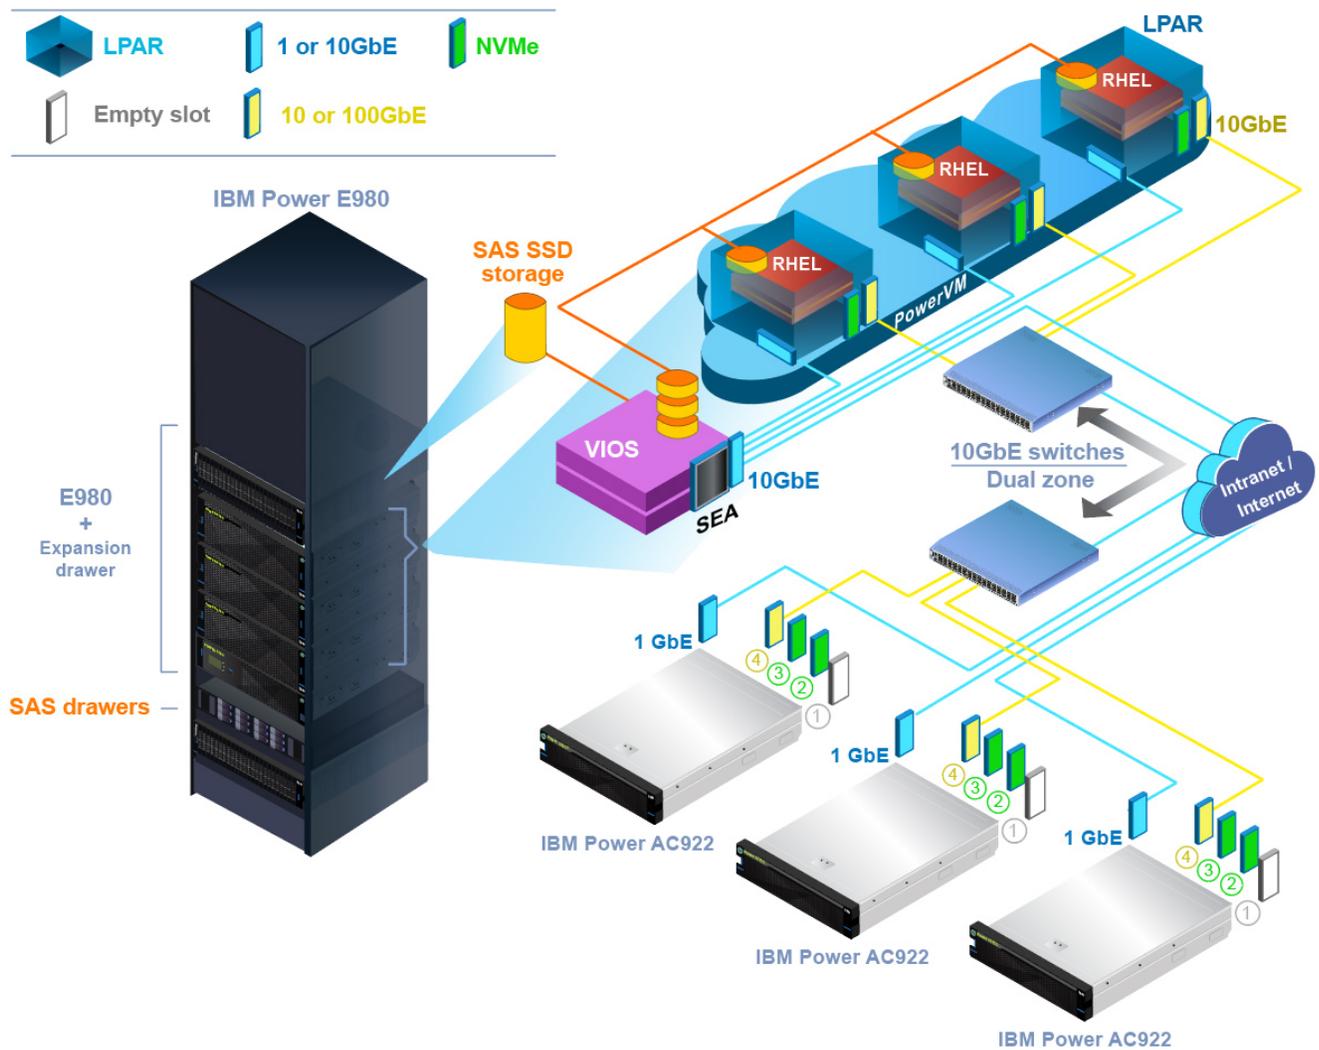


Figure 2. Architecture (MPP mixed environment)

Testing the infrastructure

Table 1 shows how the different nodes on this MPP deployment were configured. The CAS controller, SPRE, and a worker node were installed on an LPAR each on the Power E980 enterprise server and a second worker node was configured on a Power AC922 server. In this example, **sasdata** is an XFS2 local file system containing SAS data (in SAS proprietary, high-performance data format is called SASHDAT) that the Viya workload uses for running Viya jobs.

However, in general, the data to be analyzed can come in from a myriad of sources such as Network File System (NFS), file system on a storage area network (SAN) or network-attached storage (NAS), external databases, Hadoop clusters, and so on to name a few.

Mixed MPP (Power E980 LPARs / Power AC922)	Power E980 LPAR 1 (runtime environment /services)	Power E980 LPAR 2 (controller node)	Power E980 LPAR 3 (worker node 1)	Power AC922 (worker node 2)
Cores / Memory	10c / 384 GB	20c / 512 GB	30c / 1 TB	32c / 512 GB
sasdata	1 x 361 GB	2 x 361 GB	1 x 361 GB	1 x 1 TB
cascache		1 x 361 GB	2 x 361 GB	1 x 1 TB
Network	10 Gbps	10 Gbps	10 Gbps	10 Gbps

Table 1. MPP nodes

Viya uses the concept of **CAS disk cache** which acts as an extension of memory on a device. It is an on-disk repository for backing memory mapped data blocks of Viya processes. In addition to backing some data block by default, the cache gets used extensively when you run out of memory by a large margin for whatever reason (incorrect sizing, unexpected data growth, and so on). Besides this, each node in a Viya cluster replicates the changed data blocks from other nodes for fail over purposes. So, if a Viya worker node is down, the rest can keep things moving although a little slower. To achieve good performance, high-speed, low-latency, front-facing NVMe drives local to the server are recommended as done in this example.

Test results

The Viya workload consists of 29 jobs (21 Visual Statistics and 8 Visual Data Mining and Machine Learning - jobs). The generation of SASHDAT files and workload execution completed successfully in this test environment and no errors were seen in the job log files.

The Viya runtime environment reads SASHDAT format files and converts them into in-memory table representation.

It is also worth noting that the deployment and workload ran without any problems even when the deployment was spread across multiple different physical servers namely Power E980 and Power AC922. The deployment made effective use of the simplified, fast networking among LPARs on the Power E980 server where the data does not need to move over the wire unlike when it needs to communicate with the Power AC922 worker node during which time the data will need to flow over the 10GbE network as shown in Figure 2.

Summary

The successful deployment as well as workload execution concludes that this deployment is successful. It also demonstrates the superior flexibility and performance combination of the IBM Power System E980 and AC922 servers.

It reinforces effective utilization and scaling by allocating unused resources on enterprise-class scale-up Power E980 servers. This is a bonus for traditional SAS 9.4 on Power customers by giving them an option to co-exist with SAS Viya on the same server. Networking among nodes is greatly simplified and this is a critical design and sizing consideration for MPP deployments.

Get more information

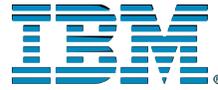
To learn more about the IBM and SAS products and capabilities, contact your IBM representative or IBM Business Partner, or visit the following websites:

- [IBM Power E980](#)
- [SAS Institute](#)
- [IBM Power Systems for Viya 3.5 Deployment Guide](#)

About the authors

Abhijit Mane is a technical consultant in IBM Cognitive Systems ISV Enablement organization. You can contact him at abhijman@in.ibm.com or <https://www.linkedin.com/in/abhijitmane>

Beth Hoffman is an IBM solution architect in IBM Cognitive Systems ISV Enablement organization. You can contact her at bethvh@us.ibm.com or www.linkedin.com/in/bethhoffmanibm



© Copyright IBM Corporation 2020
IBM Systems
3039 Cornwallis Road
RTP, NC 27709

Produced in the United States of America

IBM, the IBM logo and ibm.com are trademarks or registered trademarks of the International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked items are marked on their first occurrence in the 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

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

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

