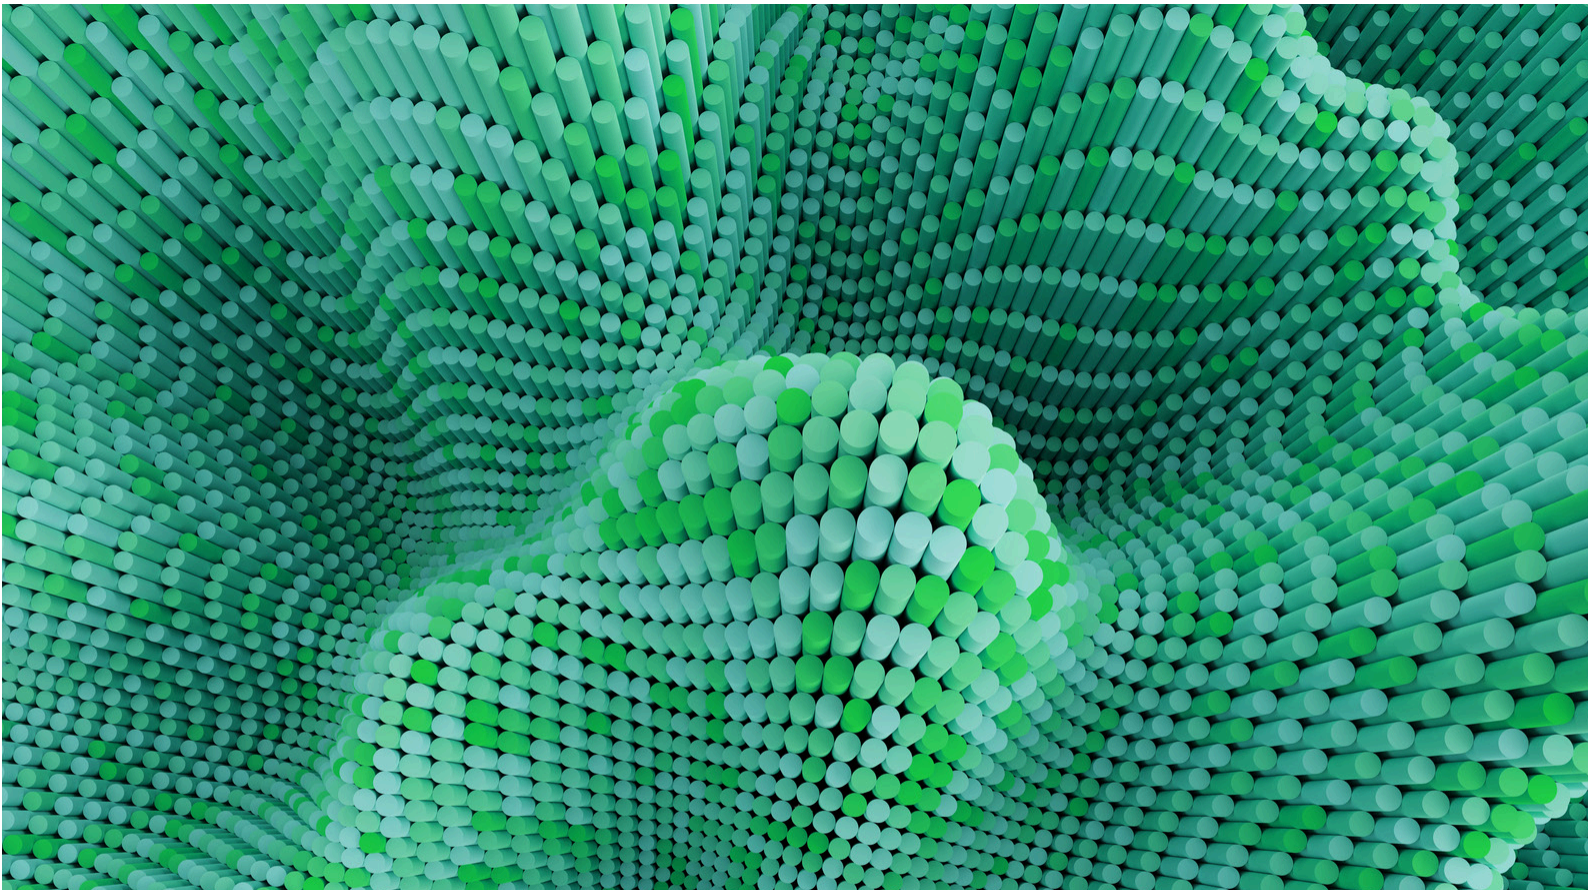
Technical white paper

SAS® Viya® on HPE and Red Hat OpenShift

HPE GreenLake Flex Solution built for SAS Viya is a validated, high-performance, on-premises, pay-as-you-go solution



Contents

Introduction ... 3
SAS Viya ... 3
The case for SAS Viya on-premises ... 3
Solution overview ... 3
SAS Viya workload classes ... 4
Deployment options ... 5
SAS Viya packages ... 5
Red Hat OpenShift ... 6
Red Hat OpenShift Virtualization ... 8
PostgreSQL database ... 8
Solution components ... 9
HPE ProLiant Gen11 servers ... 9
HPE Alletra Storage MP B10000 ... 10
HPE CSI driver ... 11
Networking ... 11
Hardware sizing and configurations ... 11
Performance ... 12
Setup ... 12
SAS Compute benchmark ... 14
SAS CAS benchmark ... 17
Summary ... 19



Introduction

In today's data-driven world, organizations across industries are racing to harness the power of analytics and artificial intelligence (AI) to gain competitive advantages, drive innovation, and optimize operations. As businesses grapple with unprecedented volumes of data, analytics emerges as the cornerstone for extracting actionable insights and driving informed decisions. Complementing analytics, AI stands at the forefront of innovation, enabling machines to learn from data, adapt to changing environments, and perform tasks that traditionally required human intelligence.

Along with massive data volumes, another challenge that enterprises face today is how distributed their data is across various data locations such as data centers, edge locations, and the public cloud. In short, their data landscape is hybrid, either by design or by accident, and it is not an easy task for them to decide whether they want to move their data entirely to the public cloud or keep it exclusively on-premises. For that reason, it would be ideal for these enterprises to have the best of both worlds: a cutting-edge AI and analytics solution that gives them the freedom of choice to keep or move their data across edge, data centers, colocation facilities, and the public cloud. This is the sweet spot that three market leaders—Hewlett Packard Enterprise, SAS, and Red Hat®—are jointly aiming for with the new HPE GreenLake Flex Solution built for SAS Viya.

This paper discusses the architecture, deployment considerations as well as performance results for this flex solution jointly created by HPE, Red Hat, and SAS.

Target audience

The audience for this paper includes chief technology officers, as well as business and IT leaders who evaluate and decide where and how their enterprise organizations should invest in the field of data analytics and AI: on-premises, hybrid, or public cloud. The paper is also intended for SAS administrators, Kubernetes administrators, and SAS users who want to understand the architectural aspects of the HPE solution for SAS Viya on Red Hat OpenShift® and begin implementing it in their environments.

SAS Viya

SAS Viya is a high-performance AI and analytics platform based on a modern, containerized architecture that delivers data-driven decisions with high speed and agility. SAS Viya revolutionizes and democratizes the entire AI and data analytics lifecycle. It can access data from virtually any data source and integrate code written in several open-source programming languages in addition to the SAS programming language—such as Python, Java, R, and Lua. It offers several UI-driven low-code or no-code tools to seamlessly manage data, develop models, and deploy insights. SAS Viya offers solutions in the areas of generative AI, machine learning (ML), deep learning, natural language processing, computer vision, neural networks, intelligent decisioning, advanced analytics, and more. It is supported on-premises as well as in the public cloud. The on-premises option is on Red Hat OpenShift and upstream Kubernetes.

The case for SAS Viya on-premises

Although the public cloud makes it relatively simple to deploy and consume IT and IT services, thousands of large enterprises clearly find it very difficult to migrate their entire IT landscapes to the public cloud. The simplest reason is the distributed hybrid nature of their data, as mentioned in the Introduction. Apart from this hybrid nature, other factors that guide a migration decision include performance, security, fragmented data sources, lack of time and resources to plan, test, and execute data consolidation, compliance, data protection policies, exceedingly high egress costs, and so on. Given all these considerations, many enterprises are now tending toward a hybrid cloud approach, deploying some workloads in the public cloud while keeping other workloads on-premises. With the on-premises HPE GreenLake Flex Solution built for SAS Viya, existing SAS®9 customers as well new ones can easily modernize to SAS Viya.

Solution overview

HPE, Red Hat, and SAS have jointly created this solution for SAS Viya on Red Hat OpenShift for on-premises or hybrid cloud environments. In this solution, SAS Viya runs on Red Hat OpenShift, which is deployed on HPE ProLiant Gen11 servers and HPE Alletra Storage MP B10000. This solution offers the choice to deploy SAS Viya on Red Hat OpenShift in either bare metal or virtual fashion. These options are described in detail later in this paper. Figure 1 shows the end-to-end solution architecture. The installation, configuration, and management of the infrastructure and the Red Hat components can either be managed by the customer themselves or be offered as a service by HPE or a partner. However, SAS Viya installation, configuration, and management are the responsibility of either the customer, a partner, or SAS, not HPE.

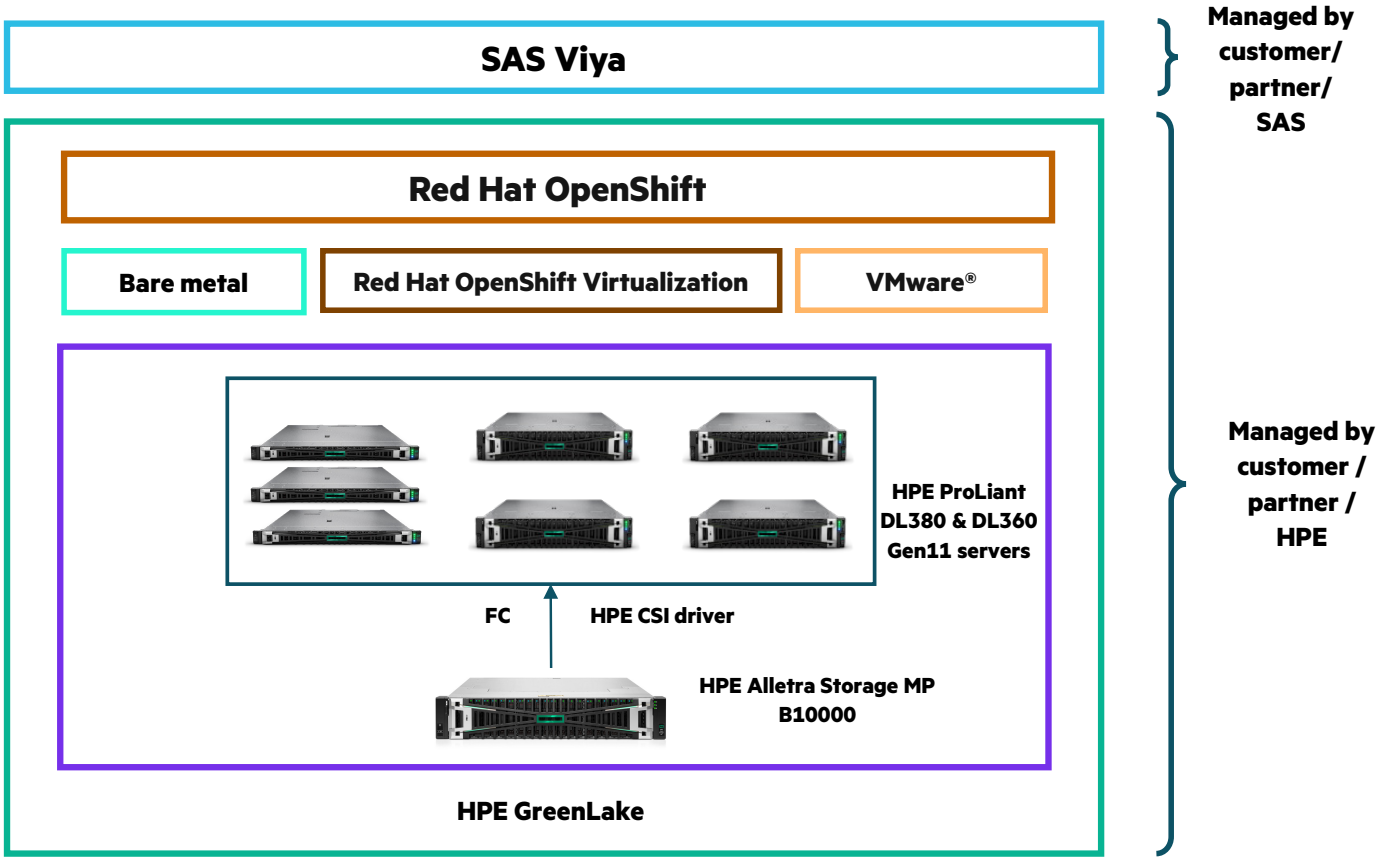


Figure 1. HPE GreenLake Flex Solution built for SAS Viya

SAS Viya workload classes

SAS Viya consists of various workload classes. Each workload class has different requirements for compute resources.

Cloud Analytic Services

Cloud Analytics Services (CAS) is the main processing engine in SAS Viya. It can perform in-memory processing of data for faster results. It requires a high CPU core count as well as a significant amount of memory. Depending on the use case, CAS can run in two modes: symmetric multiprocessing (SMP) and massively parallel processing (MPP).

CAS has two kinds of pods: CAS controller and CAS workers. The CAS controller schedules jobs on one or more CAS workers.

Compute server

Compute server processes in SAS Viya perform several functions, including but not limited to batch jobs, interactive SAS Studio sessions, and parallel compute processes that are part of SAS Model Studio.

SAS Viya services

SAS Viya has a set of stateful and stateless services. The stateless services mainly include web applications and microservices. They scale horizontally to handle increased workload and provide high availability. Stateful services, on the other hand, are open-source components that are used by SAS Viya to store important information to ensure operational integrity.

Deployment options

This flex solution built for SAS Viya offers three options for deployment: bare metal, virtualized with Red Hat OpenShift Virtualization, and virtualized with VMware.

Bare metal

In this option, the Red Hat OpenShift is deployed on bare metal HPE ProLiant Gen11 servers and HPE Alletra Storage MP B10000 *without any hypervisor*. There are two advantages of this approach: first, it is easier to deploy and second, the total cost of the solution is much lower because customers don't need to pay for the additional license costs for virtualization.

Virtualized with Red Hat OpenShift Virtualization

The first virtualized approach is to use Red Hat OpenShift Virtualization. This is a native virtualization technology from Red Hat and is being used in this solution in a hosted control planes architecture using the same servers and storage from HPE. The Red Hat OpenShift Virtualization section of this paper describes this architecture in detail.

Virtualized with VMware

The second virtualization approach is using VMware. Here, a VMware vSphere® cluster is created on the underlying HPE ProLiant Gen11 servers and HPE Alletra Storage MP B10000. Virtual machines (VM) are carved based on the sizing requirements for SAS Viya, on top of which Red Hat OpenShift is installed. Lastly, SAS Viya is deployed on this OpenShift cluster.

SAS Viya packages

SAS Viya is available in three different packages: **SAS Viya**, **SAS Viya Advanced**, and **SAS Viya Enterprise**. Table 1 shows the differences between these three packages as highlighted on the SAS Viya Capabilities and Offerings web page.

Table 1. SAS Viya packages and their differences

Features	SAS Viya	SAS Viya Advanced	SAS Viya Enterprise
Data access and in-database processing	Yes	Yes	Yes
Data preparation and integration	Yes	Yes	Yes
Data catalog, quality, and governance	Yes	Yes	Yes
Workload management	Yes	Yes	Yes
Visualization and reporting	Yes	Yes	Yes
Statistics	Yes	Yes	Yes
Matrix programming		Yes	Yes
Machine learning and deep learning	Yes	Yes	Yes
Model deployment and management	Yes	Yes	Yes
Forecasting		Yes	Yes
Text analytics		Yes	Yes
Optimization		Yes	Yes
Econometrics		Yes	Yes
Event stream analytics			Yes
Intelligent decisioning			Yes

Within each of these packages, SAS has defined small, medium, and large T-shirt sizes. These T-shirt sizes vary in factors such as the number of concurrent sessions, the number of concurrent users, the dataset size loaded in memory, and much more. As an example, Table 2 lists some of the SAS Viya Enterprise features and explains which usage assumptions were made for each T-shirt size. Similar assumptions exist for all SAS Viya packages.

Table 2. Some differences between SAS Viya Enterprise small, medium, and large T-shirt sizes

SAS Viya Enterprise features	Small	Medium	Large
Total data size loaded into disk	250 GB	1 TB	2.5 TB
Total data lifted into memory	100 GB	600 GB	1.5 TB
SAS Visual Statistics using linear regression	Concurrent sessions = 1, Dataset size = 2 GB	Concurrent sessions = 2, Dataset size = 5 GB	Concurrent sessions = 2, Dataset size = 10 GB
SAS Model Manager: number of models stored in model repository	10	20	50
SAS Visual Forecasting: maximum number of time series to be forecasted in a time series table	25,500	150,000	300,000

As the features grow from small to medium to large in SAS Viya, so do their hardware requirements. This is the basis for defining hardware configurations within this flex solution for SAS Viya. These configurations are described in the Hardware sizing and configurations section of this paper.

Red Hat OpenShift

Red Hat OpenShift brings together tested and trusted services to reduce the friction of developing, modernizing, deploying, running, and managing applications. OpenShift delivers a consistent experience across public cloud, on-premises, hybrid cloud, or edge architecture.

Red Hat OpenShift offers a unified, flexible platform to address a variety of business needs spanning from an enterprise-ready Kubernetes orchestrator to a comprehensive cloud-native application development platform that can be self-managed or used as a fully managed cloud service. Red Hat OpenShift provides:

- The ability to deploy and run in any environment, the flexibility to build new applications, modernize existing applications, run third-party ISV applications, or use public cloud services under a single platform
- The tools needed to help customers integrate data analytics and AI/ML capabilities into cloud-native applications to deliver more insight and value
- Consistency and portability to deploy and manage containerized workloads, make infrastructure and investments future-ready, and deliver speed and flexibility on-premises, across cloud environments, and to the edge of the network
- Advanced security and compliance capabilities, enabling end-to-end management and observability across an entire architecture

Red Hat Advanced Cluster Security for Kubernetes modernizes container and Kubernetes security, enabling developers to add security controls early in the software lifecycle. Red Hat Advanced Cluster Management for Kubernetes lets you manage your entire application lifecycle and deploy applications on specific clusters based on labels, and Red Hat OpenShift Data Foundation supports performance at scale for data-intensive workloads.

In addition, the platform delivers a suite of tools to help teams build and scale applications, including OpenShift Serverless, OpenShift Service Mesh, GitOps, and OpenShift Pipelines to deploy modern architecture and operational practices such as serverless application programming interfaces (APIs), event-driven architecture, and automated pipelines to simplify application development, delivery, and integration. Using these technologies, developers can focus on creating applications that deliver business value—without requiring detailed knowledge of the underlying infrastructure. Applications can be developed, and their lifecycle managed at scale with more security, more consistency, and more reliability across a wide variety of systems.

Red Hat OpenShift also extends the capabilities of Kubernetes to the network edge, helping organizations manage infrastructure consistently at scale, even to the most remote edge locations, without sacrificing security or stability. As part of a hybrid cloud environment, it extends processing closer to users and data sources.

Deploying SAS Viya on OpenShift brings many advantages to the overall deployment, including these:

- **Improved security posture:** As applications are deployed or updated, it is critical to provide dynamic security controls to help keep the business as safe as possible. OpenShift can apply security controls to the software supply chain, improving the security capabilities of applications without reducing developer productivity. Red Hat OpenShift enables businesses to control, defend, and extend their application platform throughout its lifecycle.
  - **Governance and compliance:** Red Hat Advanced Cluster Management for Kubernetes, included in Red Hat OpenShift Platform Plus, can help organizations more quickly deploy apps, manage multiple clusters, and enforce policies across clusters at scale. OpenShift offers core security capabilities such as access controls, networking, and enterprise registry with a built-in scanner. Red Hat Advanced Cluster Security for Kubernetes enhances this advantage with additional security capabilities, including runtime threat detection, full lifecycle vulnerability management, and risk profiling.

- **Manageability and flexibility:** OpenShift simplifies the deployment and management of a hybrid infrastructure, giving you the flexibility to have a self-managed or fully managed service running on-premises or in cloud and hybrid environments. Red Hat Advanced Cluster Management controls clusters and applications from a single console, with built-in security policies.

Control plane nodes

A node is a virtual or bare metal machine in a Kubernetes cluster. Control plane nodes host the Kubernetes services that are needed to control the Kubernetes cluster, such as the Kubernetes API server, etcd, Kubernetes controller manager, and the Kubernetes scheduler. In OpenShift, the control plane nodes run additional services, including the OpenShift API server, OpenShift controller manager, OpenShift OAuth API server, and OpenShift OAuth server.

By default, the OpenShift installer deploys a highly available OpenShift control plane composed of three control plane nodes. Additional options are available for running the control plane in standalone mode for smaller environments, such as installing onto a single node or compact three-node cluster where the control plane nodes are marked as schedulable, enabling them to also act as worker nodes to run end-user applications.

- A management cluster may also be used to consolidate multiple control planes in hosted mode for multiple clusters using Red Hat hosted control planes (HCP), which is based on the open-source project HyperShift. Using HCP with OpenShift provides several benefits, including these:

- **Greater security** by decoupling the control plane from the data plane
- **Lower costs** by running many control planes on fewer nodes
- **Faster startup times** for control planes because the control planes consist of pods that are launched on OpenShift

Worker nodes

Worker nodes host the end-user application container pods. OpenShift environments can have many worker nodes. However, a minimum of two worker nodes is used to provide node redundancy. The total count and sizing of worker nodes is determined by the application resource requirements plus additional capacity to plan for node maintenance operations, such as cluster updates.

SAS Viya workload classes and node pools can be used to place the workloads on appropriate worker nodes within the cluster. The placement of SAS workload classes is enabled by applying predefined Kubernetes node labels and node taints. SAS Viya workload placement using node pools is not mandatory, but it does provide many benefits, such as enabling you to tailor the cluster topology to workload requirements, for example, by specifying different hardware configurations for nodes with additional storage, with GPU cards, and so forth.



Red Hat OpenShift Virtualization

OpenShift Virtualization enables you to run and manage VM workloads alongside container workloads. OpenShift Virtualization gives you the capability to deploy SAS Viya on a tenant OpenShift cluster running on VMs as part of this flex solution built for SAS Viya. Hewlett Packard Enterprise recommends that you leave hyperthreading on in OpenShift Virtualization to have two logical CPUs on each physical core.

Figure 2 shows the components for the solution, which uses HCP with OpenShift Virtualization within a management cluster that runs on bare metal HPE ProLiant servers. This supports the installation of multiple-tenant OpenShift clusters that utilize worker nodes running on VMs.

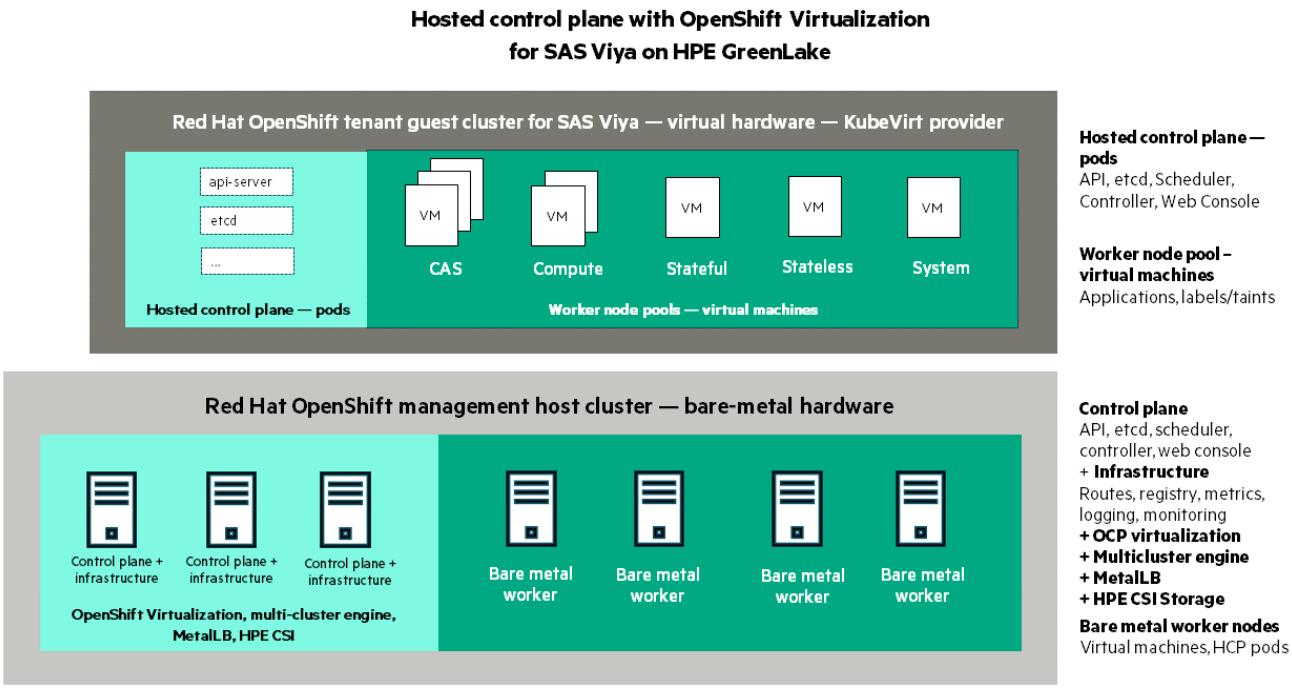


Figure 2. Red Hat OpenShift HCP with OpenShift Virtualization

Red Hat and Intel® have worked extensively to enhance the support for Red Hat technologies. For specific AI and ML use cases, patterns validated by Red Hat take advantage of specific hardware accelerators built into the 4th Gen Intel® Xeon® Scalable processors, such as Intel Advanced Matrix Extensions (Intel AMX), purpose-built to accelerate AI inference and training.

SAS is validated to run on the 5th Gen and 4th Gen Intel Xeon Scalable processors.

PostgreSQL database

SAS Viya requires a PostgreSQL database for the SAS Infrastructure Data Server to store SAS Viya platform user content, such as reports, authorization rules, selected source definitions, attachments, and audit records. The SAS Viya Platform Operations document describes resource requirements for the PostgreSQL server in detail. During the deployment, customers can choose to either use the PostgreSQL database shipped with the software or deploy a separate distribution of PostgreSQL externally to SAS Viya. This PostgreSQL server instance can be hosted on either a bare metal server or a virtual server outside the OpenShift cluster or on the same OpenShift cluster as SAS Viya but in a separate namespace. The T-shirt sizing described in the Hardware sizing and configurations section of this paper assumes that the PostgreSQL database will be deployed on the same OpenShift cluster outside the SAS Viya namespace if a customer opts for using an external PostgreSQL database.

Solution components

This section describes the compute, storage, and networking components of the HPE GreenLake Flex Solution built for SAS Viya.

HPE ProLiant Gen11 servers

Featuring the latest innovations and engineered for hybrid environments, HPE ProLiant Gen11 servers are built around three design principles to deliver:

- An intuitive cloud operating experience
- Trusted security by design
- Optimized performance

For business-critical workloads such as SAS Viya, where speed and agility are critical, HPE ProLiant DL3x0 rack mount servers are optimally suited to fulfill compute requirements to scale quickly and efficiently. The HPE ProLiant Gen11 servers in this flex solution built for SAS Viya include the following notable characteristics:

- 5th Gen Intel Xeon Scalable processors (Emerald Rapids)
- Eight channels of DDR5 memory at up to 5600 MT/s (max 4 TB/socket).
- PCIe Gen5 slots providing support for peripherals such as GPU accelerators, storage, and networking cards—to cite a few—at double the bandwidth of the previous generation (up to 32 GT/s)

In addition to the improvements of more cores and faster memory, the 5th Gen Intel Xeon Scalable processors offer several new CPU features and innovative accelerators:

- Significantly larger loadline calibration (LLC): 3x larger on top-bin 5th Gen compared to 4th Gen
- Trust Domain Extensions (TDX) support
- Improvements in idle power and power efficiency under load

And they leverage the following features introduced in the 4th Gen:

- Intel Advanced Matrix Extensions (Intel AMX) to speed up low-precision math and accelerate AI/ML
- Intel Data Streaming Accelerator (Intel DSA) to copy and move data faster and assist PDK
- Intel QuickAssist Technology (Intel QAT) to accelerate compression, encryption, and decryption
- Intel In-memory Analytics Accelerator (Intel IAA) to speed up query processing performance
- Intel Dynamic Load Balancer (Intel DLB) to help speed up data queues

HPE ProLiant is engineered with a fundamental security approach to defend against increasingly complex threats through an uncompromising focus on state-of-the-art security innovations that are built into the DNA of Hewlett Packard Enterprise—from silicon to software, from the factory to the cloud. Silicon root of trust from HPE, enabled by HPE iLO, secures millions of lines of firmware code and more than four million HPE servers around the world.

HPE ProLiant Gen11 offers enhanced operational efficiency compared to previous generations through the modernization of workload infrastructure. A remarkable 8:1 consolidation ratio has been demonstrated when transitioning from HPE ProLiant Gen9 to HPE ProLiant Gen11 servers, and an impressive 11:1 consolidation ratio when upgrading from HPE ProLiant Gen8 to Gen11 servers. This strategic server consolidation not only optimizes performance but also streamlines resource utilization, thereby improving sustainability, thus aligning your IT infrastructure with the latest advancements in technology.



The HPE GreenLake Flex Solution built for SAS Viya comprises two different HPE ProLiant Gen11 server models, as shown in Figure 3. These models cater to resource requirements for different SAS Viya packages and T-shirt sizes. The HPE hardware configurations for this solution are explained in detail in the Hardware sizing and configurations section of this paper. What is important to note is that the 5th Gen Intel Xeon Scalable processor model used in these servers offers high CPU frequency and optimum number of cores required to run a wide range of SAS Viya use cases with demanding performance requirements.



HPE ProLiant DL380 Gen11
2U, 2P

HPE ProLiant DL360 Gen11
1U, 2P

Figure 3. HPE ProLiant DL Gen11 servers for SAS Viya

Note
The hardware configurations in this flex solution do not explicitly include GPUs. However, the HPE ProLiant DL380 Gen11 server supports a variety of NVIDIA® and Intel GPUs. If the SAS Viya use case for a particular customer requires GPUs, they can be added on to the bill of material. For more information, refer to the HPE ProLiant DL380 Gen11 QuickSpecs.

HPE Alletra Storage MP B10000

The storage component of this solution is HPE Alletra Storage MP B10000. It is an AI-driven, all-NVMe, cloud-native, and hybrid-by- design block storage solution. With a highly flexible and scalable architecture, HPE Alletra Storage MP B10000 leaps far ahead of the traditional monolithic arrays. It is the industry's first disaggregated, scale-out block storage system, which means you can independently scale performance and capacity, in other words, compute and storage, depending on your requirements. It offers extreme performance, resiliency, and efficiency for all kinds of workloads and applications.

With a 100% availability guarantee[1] and the ability to seamlessly consolidate multiple workloads, HPE Alletra Storage MP B10000, shown in Figure 4, helps to ensure that your mission-critical and data-intensive workloads can continue to run without disruptions. Powered by HPE GreenLake, it offers a cloud-operational experience along with cloud-like consumption. As a customer, you are free to choose between traditional CapEx pricing or consumption-based pay-per-use[2] pricing.



Figure 4. HPE Alletra Storage MP B10000

[1] Uptime guarantees that deliver, April 2023
[2] May be subject to minimums or capacity may apply reserve

HPE CSI driver

To access storage resources and perform efficient data management operations from Red Hat OpenShift, HPE provides a Container Storage Interface (CSI) driver. This flex solution created for SAS Viya always uses the latest version of the HPE CSI driver available in the Red Hat OpenShift Operator Hub. This CSI driver has certified HPE Alletra Storage MP B10000 with Red Hat OpenShift, which is used in this flex solution. More information about HPE CSI driver is available at HPE CSI Driver for Red Hat OpenShift.

As mentioned in the HPE CSI driver documentation, while configuring the HPE CSI driver for SAS Viya, add the fsMode: "0770" parameter to the StorageClass with RWO claims or fsMode: "0777" for RWX claims.

Networking

This flex solution uses Fibre Channel as the protocol to connect the underlying HPE Alletra Storage MP B10000 to the HPE ProLiant Gen11 servers. In the XSMALL and SMALL configurations, we have 2x HPE Storage Fibre Channel Switch C-series SN6610C. For the larger configurations, we have 2x HPE Storage Fibre Channel Switch C-Series SN6720C.

Hardware sizing and configurations

This section describes hardware configurations and sizing for the HPE GreenLake Flex Solution built for SAS Viya. SAS has published hardware estimation guides for SAS Viya T-shirt sizes available across three SAS Viya packages. These guides describe various SAS Viya features, the total data size residing on disk, the total data size loaded in memory, and the number of concurrent users and sessions. Most importantly, they also describe the hardware resource estimate required for each SAS Viya workload class. HPE and Red Hat followed these hardware estimates and created the corresponding hardware configurations for both bare metal and virtualized deployments.

In creating these hardware configurations, HPE has taken into consideration factors such as efficiency, performance requirements, total cost of the solution, and ease of deployment. For example, the extra small bare metal (HPE_SASVIYA_BM_XSMALL_B_FC) and virtualized (HPE_SASVIYA_V_XSMALL_B_FC) configurations are intended for cost-sensitive and low-performance deployments that might be appropriate for customers who want to start with a small SAS Viya environment. Similarly, the medium and large configurations are for the high-scale performance-oriented SAS Viya deployments. Care has been taken to judiciously include resources to avoid over-provisioning while at the same time leaving enough headroom for SAS Viya node pools to grow.

In this flex solution created for SAS Viya, all the storage needed for SAS Viya, including persistent storage, ephemeral storage, and so on, is served by the HPE Alletra Storage MP B10000.

These hardware configurations are available in the HPE One Config Advanced tool. They serve as starting points for a customer opportunity and can be customized based on specific customer use cases and requirements.

Table 3 describes the HPE hardware configurations for a bare metal deployment with no hypervisor. Across all the configurations, three HPE ProLiant DL360 Gen11 servers, with a lower-end Silver series 5th Gen Intel Xeon Scalable processor form the Red Hat OpenShift control plane. For worker nodes, we have two HPE ProLiant DL360 Gen11 servers with Gold series processors as the worker nodes in the XSMALL and SMALL configs. Note that for these two configs, the storage system is the same HPE Alletra Storage MP B10000 8c 2N but with varying numbers of NVMe SSDs. When it comes to MEDIUM, LARGE and XLARGE configs, we have three to six HPE ProLiant DL380 Gen11 servers with Gold series processors. The underlying storage is the HPE Alletra Storage MP B10000 from the 16-core 2 nodes to the larger 16-core 4 nodes variant. This sizing has taken into consideration results from the performance testing done by HPE and SAS for this flex solution as described in the Performance section of this paper.

Table 3. HPE GreenLake Flex Solution built for SAS Viya bare metal configurations

HPE bare metal configs for SAS Viya	HPE_SASVIYA_BM_ XSMALL_B_FC	HPE_SASVIYA_BM_ SMALL_B_FC	HPE_SASVIYA_BM_ MED_B_FC	HPE_SASVIYA_BM_ LARGE_B_FC	HPE_SASVIYA_BM_ XLARGE_B_FC
SAS Viya T-shirt	SAS Viya Small	SAS Viya Advanced Small SAS Viya Enterprise Small	SAS Viya Medium SAS Viya Advanced Medium	SAS Viya Enterprise Medium SAS Viya Large	SAS Viya Advanced Large SAS Viya Enterprise Large
Platform	Red Hat OpenShift 4.15	Red Hat OpenShift 4.15	Red Hat OpenShift 4.15	Red Hat OpenShift 4.15	Red Hat OpenShift 4.15
Physical servers	3x HPE ProLiant DL360 Gen11 & 2x HPE ProLiant DL360 Gen11	3x HPE ProLiant DL360 Gen11 & 2x HPE ProLiant DL360 Gen11	3x HPE ProLiant DL360 Gen11 & 3x HPE ProLiant DL380 Gen11	3x HPE ProLiant DL360 Gen11 & 4x HPE ProLiant DL380 Gen11	3x HPE ProLiant DL360 Gen11 & 6x HPE ProLiant DL380 Gen11
Storage array	HPE Alletra Storage MP B10120	HPE Alletra Storage MP B10120	HPE Alletra Storage MP B10130	HPE Alletra Storage MP B10140	HPE Alletra Storage MP B10230

Table 4 shows the virtualized HPE configurations for this flex solution. If SAS Viya is being deployed using either OpenShift Virtualization or VMware, these hardware configurations must be used.

Note that in the case of virtualized configurations, the underlying HPE ProLiant DL360 or DL380 Gen11 servers share resources to host the Red Hat OpenShift control plane as well as the VMs that act as worker nodes for running SAS Viya.

As far as the storage is concerned, the sizing is similar to the bare metal configuration because the performance requirements remain the same. The HPE Alletra Storage MP B10000 is used starting from the 8-core 2-node system to the 16-core 4-node variant. The suffix B_FC in the names of both the bare metal and virtualized configs indicate block storage and Fibre Channel.

Table 4. HPE GreenLake Flex Solution built for SAS Viya virtualized configurations

HPE virtualized configs for SAS Viya	HPE_SASVIYA_V_ XSMALL_B_FC	HPE_SASVIYA_V_ SMALL_B_FC	HPE_SASVIYA_V_ MED_B_FC	HPE_SASVIYA_V_ LARGE_B_FC	HPE_SASVIYA_V_ XLARGE_B_FC
SAS Viya T-shirt	SAS Viya Small	SAS Viya Advanced Small SAS Viya Enterprise Small	SAS Viya Medium SAS Viya Advanced Medium	SAS Viya Enterprise Medium SAS Viya Large	SAS Viya Advanced Large SAS Viya Enterprise Large
Platform	Red Hat OpenShift 4.15	Red Hat OpenShift 4.15	Red Hat OpenShift 4.15	Red Hat OpenShift 4.15	Red Hat OpenShift 4.15
Physical servers	3x HPE ProLiant DL360 Gen11	3x HPE ProLiant DL360 Gen11	4x HPE ProLiant DL380 Gen11	5x HPE ProLiant DL380 Gen11	7x HPE ProLiant DL380 Gen11
Storage array	HPE Alletra Storage MP B10120	HPE Alletra Storage MP B10120	HPE Alletra Storage MP B10130	HPE Alletra Storage MP B10140	HPE Alletra Storage MP B10230

Performance

HPE, SAS and Red Hat have worked jointly on running extensive benchmarks on the HPE GreenLake Flex Solution built for SAS Viya. We used the HPE_SASVIYA_BM_LARGE_B_FC bare metal hardware configuration to emulate a customer environment to understand the behavior and performance of the solution, and to identify deployment best practices. Here are the details of the setup and the SAS workloads that were used for testing and the results.

Setup

The performance setup had the following components:

- 3x HPE ProLiant DL360 Gen11 servers with Silver series 4th Gen Intel Xeon processors as the Red Hat OpenShift control plane
- 4x HPE ProLiant DL380 Gen11 servers with Gold series 4th Gen Intel Xeon processors as Red Hat OpenShift worker nodes
- HPE Alletra Storage MP B10000 32c 2N system
- HPE CSI driver 2.4.2
- 2x HPE Storage Fibre Channel Switch B-series SN6700B with 24 active ports

Please note that although for this performance testing the available hardware had 4th Gen Intel Xeon processors, the results apply to the hardware configurations in the flex solution mentioned in this paper because they come with the next-generation Intel processors.

Deployment considerations

Deploying SAS Viya on the HPE GreenLake cloud is not different than any other deployment of SAS Viya on the Red Hat OpenShift. Detailed information about the required steps can be found in the SAS Viya Operations Guide and in these Red Hat blog posts (part 1 and part 2).

While planning the deployment, special care was given to ensure that the available cluster resources (in terms of storage, CPU and memory) would be optimally utilized by the SAS Viya environment. Figure 5 shows a schematic of the setup for this performance testing. Note that this is one of the ways to deploy SAS Viya. The SAS Viya Platform Operations Guide describes all the supported deployment options in detail.

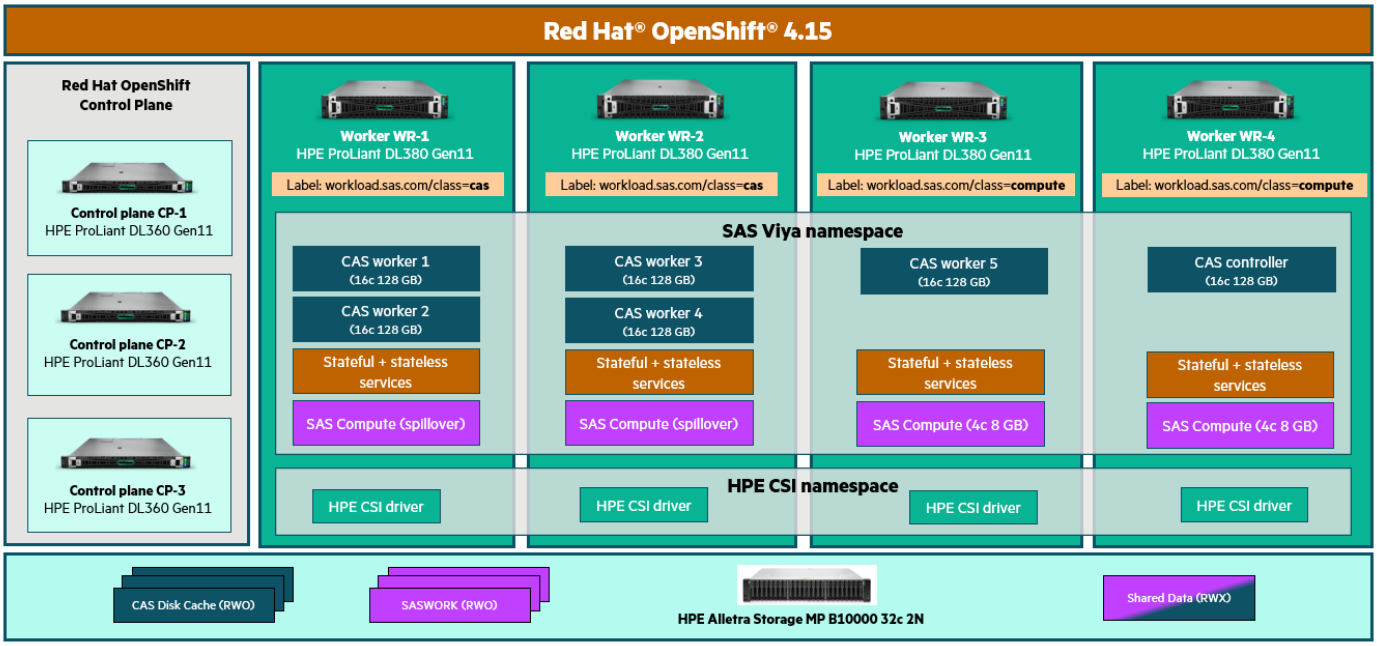


Figure 5. Performance testing setup with the HPE GreenLake Flex Solution built for SAS Viya

Storage considerations

The HPE Alletra Storage MP B10000 exposes both ReadWriteOnce (RWO) as well as ReadWriteMany (RWX) or Network File System (NFS) storage to the OpenShift cluster through the HPE CSI driver. Both storage types are needed for SAS Viya. It is a typical requirement that data sources residing on file shares need to be accessed through NFS while the compute engines in SAS Viya require fast, ephemeral storage for working efficiently with the data after ingestion (for example, for SASWORK and CAS Disk Cache).

To comply with these requirements, we provisioned a single large volume to keep the benchmark source data, which was mounted to all SAS Compute and Cloud Analytic Services (CAS) sessions as a RWX (NFS) file share. For providing the compute engines with ephemeral storage, a generic ephemeral volume configuration was used, which lets Kubernetes create RWO volumes on demand when SAS sessions are launched.

CPU and memory considerations

SAS benchmarks utilize both the major compute engines available in the SAS Viya platform: SAS Compute, which is the classic SAS kernel for both data processing and analytics, and CAS, which is the distributed in-memory server for analytical workloads.

SAS Compute sessions in SAS Viya are launched as Kubernetes jobs, which means that resource allocation is declaratively managed by a "resources" section in their template, where a start request and a maximum setting for CPU and memory can be given per session. Since many SAS procedures are multi-threaded, finding a reasonable configuration is crucial for optimizing the performance of SAS sessions. Table 5 shows the resources that were assigned to these SAS Compute sessions.

Table 5. CPU and memory assignment for SAS Compute sessions

Resource	Request ("start with")	Limit ("can use up to")
CPU	500 millicores	4
Memory	1 GB	8 GB

Note that these values heavily depend on the characteristics of your SAS Compute workload and the settings shown here might not be optimal for every case.

While SAS Compute sessions are launched as jobs and will terminate after their task has completed, the SAS CAS server is deployed as a service, in other words, CAS is expected to be up and running continuously. Following the SAS recommendations for the SAS Viya Enterprise medium T-shirt, this setup used a distributed configuration for the CAS server (often referred to as "CAS MPP"), which consists of one CAS controller and five CAS worker pods, deployed with a QoS class of "guaranteed" to prevent eventual pod evictions. Table 6 shows the CPU and memory assigned to the CAS pods.

Table 6. CPU and memory assigned to SAS CAS pods

Role	Pod count	CPU	Memory
CAS Controller	1	16	128 GB
CAS Worker	5	16	128 GB

Workload placement

Using node labels (and optionally node taints), Kubernetes administrators can change the default pod scheduling on a cluster. To optimize the usage of the available hardware resources for the benchmarks, we have only used node labels, but not taints, which gives the Kubernetes pod scheduler more flexibility. The SAS Viya Platform Operations Guide has an extensive section about workload placement and customers should carefully review it and adapt the deployment to their requirements.

As can be seen in Table 7, the intention is to use the more powerful nodes (Worker 1 and Worker 2) predominantly for the CAS server, while the remaining nodes are set as preferred targets for SAS Compute sessions. Note that setting the label workload.sas.com/class=compute is required when intending to use SAS Workload Management, which was not the case in this benchmark. Also, it is worth noting that due to the absence of node taints, this is a "soft" setting, which can be overruled by the Kubernetes pod scheduler if needed.

Table 7. Node labels in the setup for performance testing

Worker nodes	Logical CPU	Memory	Role	Node label
WR-1	64	512 GB	CAS	workload.sas.com/class=cas
WR-2	64	512 GB	CAS	workload.sas.com/class=cas
WR- 3	64	256 GB	Compute	workload.sas.com/class=compute
WR- 4	64	256 GB	Compute	workload.sas.com/class=compute

By default, the pod definitions for the SAS CAS server include an anti-affinity rule, which would prevent two CAS server pods from being deployed on the same Kubernetes node. Since this cluster in the HPE GreenLake Flex Solution built for SAS Viya cluster has fewer physical servers as worker nodes (4) than the number of CAS server pods (1+5), the anti-affinity rule needs to be removed from the pod definition. It is important to highlight that all worker nodes of the cluster in this flex solution for SAS Viya come with the same Intel Xeon chipset, so each CAS pod can work under identical conditions.

SAS Compute benchmark

The benchmark to test the SAS Compute performance is the SAS mixed analytics workload. It consists of a collection of SAS programs covering a variety of typical data processing tasks as well as analytical tasks which do not use the CAS server. Grafana dashboards provided by SAS and HPE were used to track resource usage while the tests were executed. Here are some characteristics of this benchmark:

- Data processing: SAS data steps, SQL queries, ingesting data from raw formats, sorting data, among other tasks
- Analytical: logistic regressions, linear and nonlinear mixed models, among other tasks
- A total of 306 SAS sessions (jobs) launched concurrently with only a minimal ramp-up time (a few seconds)
- More than 1.5 TB of data ingested and processed.



Results

As can be seen in Figure 6, the distribution of SAS Compute sessions clearly followed the expected pattern, given the node label settings listed in Table 7 above.

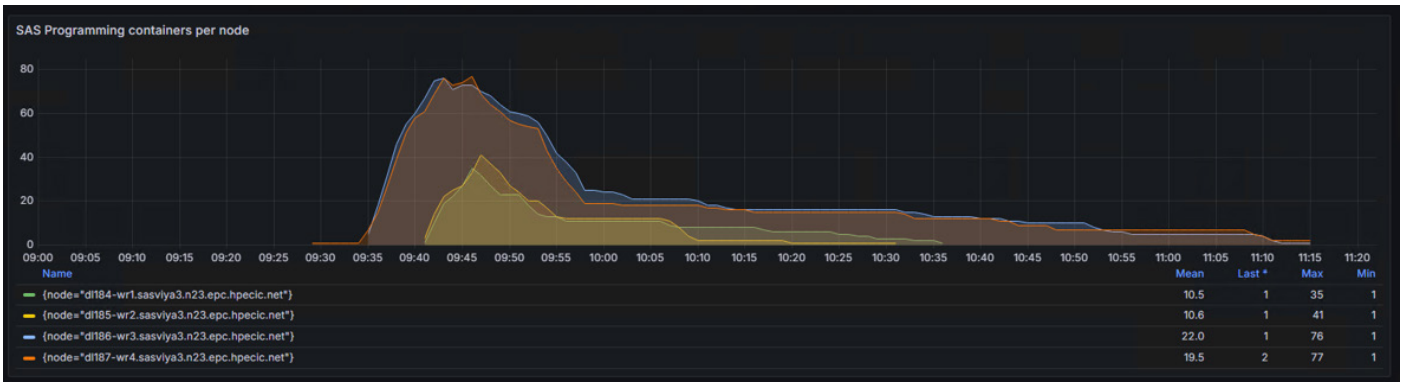


Figure 6. SAS Compute sessions per worker node

The Kubernetes scheduler preferred launching the SAS Compute sessions on nodes WR-3 and WR-4 but was able to "spill-over" new requests to nodes WR-1 and WR-2 once the other servers were saturated from a CPU perspective (at around 9:40h). This means that the average time for pods being stuck in "Pending" state was minimized.

Figure 7 indicates that at peak times, more than 220 SAS Compute sessions were running on the cluster concurrently.

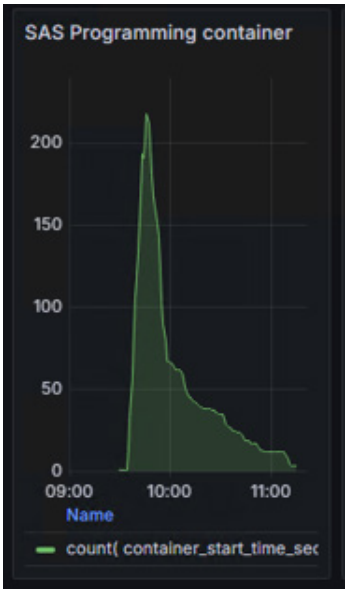


Figure 7. Number of pods running concurrently

Figures 8 and 9 indicate that the selected CPU and memory settings were sufficient for the benchmark, since the sessions rarely allocated these resources to their allowed maximum.



Figure 8. CPU usage by SAS Compute sessions



Figure 9. Memory usage by SAS Compute sessions

Figure 10 shows the storage throughput while the SAS Compute jobs were running. This graph was taken at the level of the HPE Alletra Storage MP B10000. The total system bandwidth shown here includes the I/O traffic generated by all concurrent SAS Compute sessions on both the SASWORK and the NFS file share. The graph shows a typical usage pattern common to most SAS programs which initially ingest the input data from a shared location and then work on that data in their respective SASWORK folders. At peak we saw an impressive maximum of 21 GB/s, which resolves to 168 MB/s per core used by SAS Compute. This easily surpasses the minimum I/O throughput of 100-150 MB/s per core recommended by SAS.

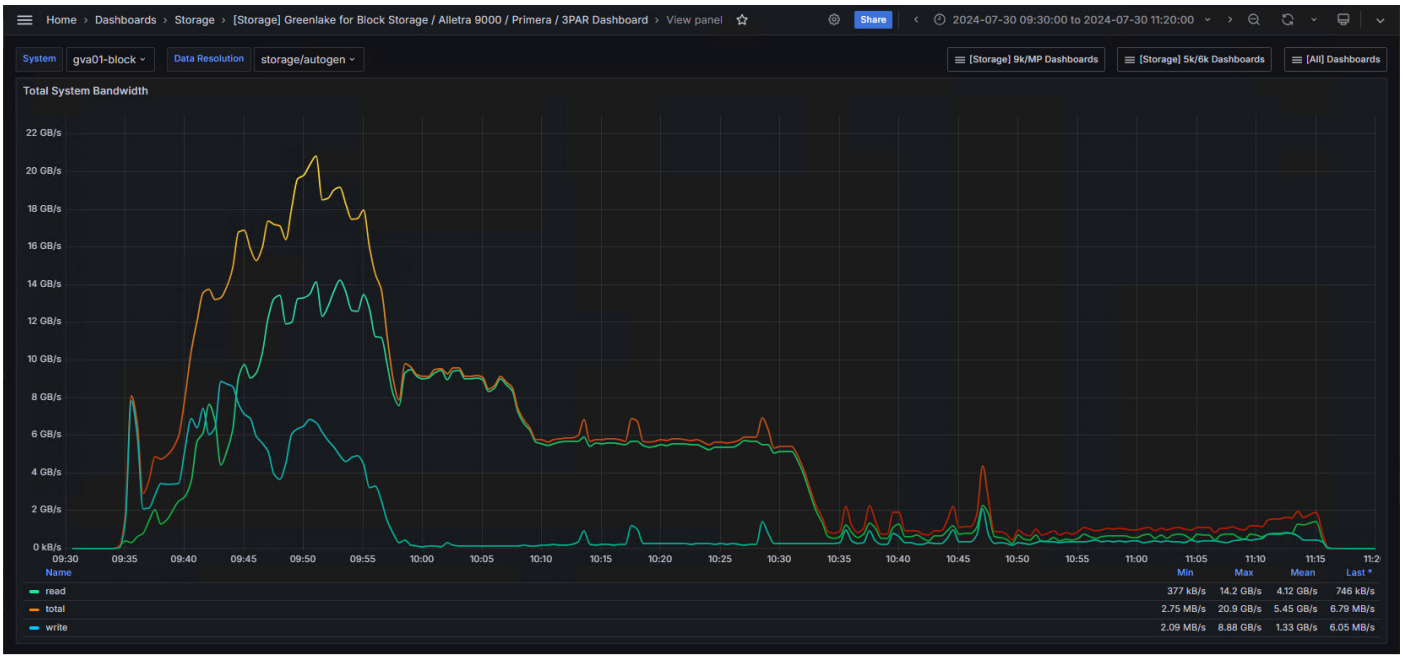


Figure 10. Combined read + write throughput on the HPE Alletra Storage MP B10000 with 306 SAS Compute sessions

SAS CAS benchmark

The benchmark for the SAS CAS Server is called the Data Mining and Machine Learning (DMML) workload. This benchmark processes about 250 GB of in-memory data using various AI and analytical procedures, including:

- Neural network
- Gradient boosting algorithms
- Decision trees
- Support vector machines

The benchmark starts with uploading the data to the CAS server and ends with writing back the result datasets to disk.

Results

The DMML benchmark was tested using different CPU and memory configurations. In all these tests, the CAS server was consistent in evenly distributing the workload and was able to fully utilize each CAS worker's assigned CPU. Figure 11 shows the CPU utilization where each CAS worker pod was assigned 16 logical cores.



Figure 11. CPU utilization by CAS workers

Even though the CAS server keeps its data in memory, it does not mean that the data resides in physical RAM. Rather, the CAS server leverages virtual memory by default, where data is partially swapped out to the disk. In general, all in-memory data is memory-mapped to a locally accessible backing store, such as the CAS Disk Cache on the underlying block storage.



Figure 12. Resident set size memory for CAS

The resident set size (RSS) memory metric clearly shows that most of the 230 GB of in-memory data was not actually residing in physical RAM throughout the benchmark. Likewise, the total system bandwidth chart in Figure 13 indicates that the system was initially busy with uploading the source data and then immediately swapping it out to the disk-based cache. The peak throughput of almost 8.5 GB/s was less than the 22 GB/s we saw in the SAS Compute benchmark, but it is important to understand that this bandwidth was consumed by just five CAS worker pods (instead of 306). This results in an MB/s rate of 217 MB/s per core.

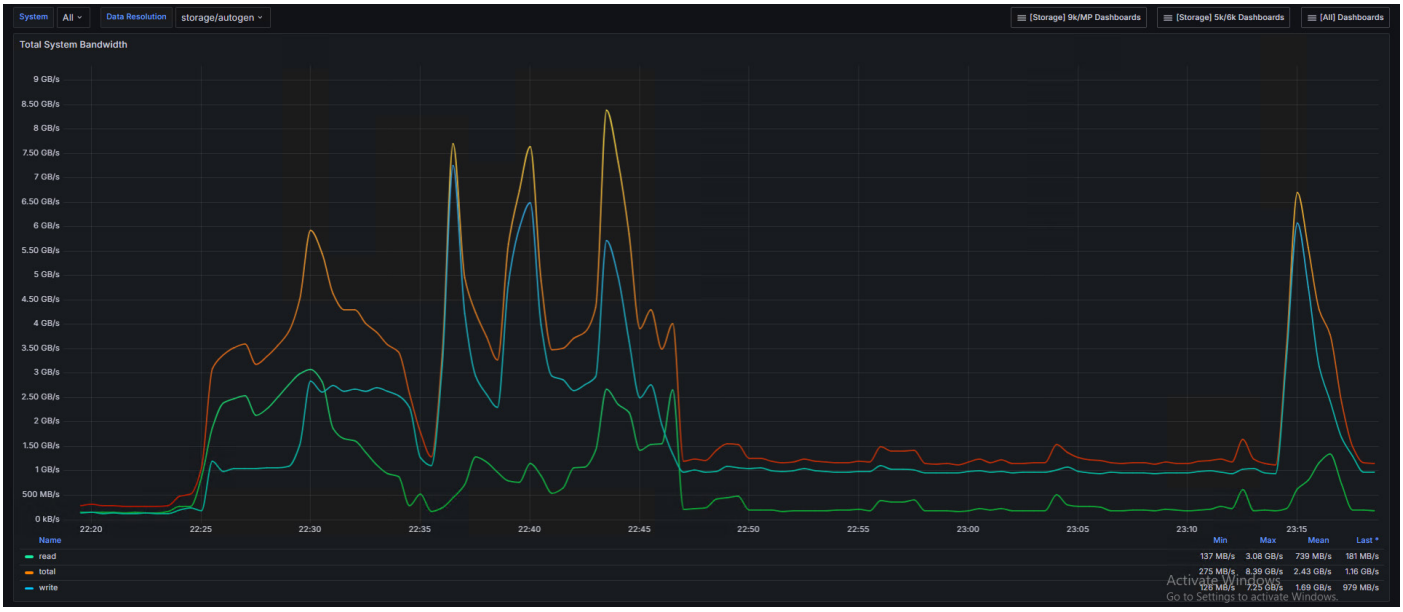


Figure 13. Combined read + write throughput for the CAS benchmark at the storage layer

Summary

The HPE GreenLake Flex Solution built for SAS Viya delivers a proven, high-performance and efficient on-premises offering for your AI and analytics use cases. This solution offers you the choice to deploy SAS Viya with Red Hat OpenShift on bare metal servers or in a virtualized fashion using OpenShift Virtualization or VMware.

The hardware components in this flex solution are HPE ProLiant Gen11 servers with 5th Gen Intel Xeon Scalable processors and HPE Alletra Storage MP B10000. HPE GreenLake provides the freedom to keep your data wherever it is and bring the cloud experience and cloud-like consumption to it instead of migrating the data to the public cloud. The HPE GreenLake enables configuration, management, and monitoring of your entire IT landscape from a single pane of glass. HPE offers managed services for the infrastructure as well as the platform including the Red Hat components. SAS Viya management is not offered by HPE for this flex solution.

Together with HPE, SAS, and Red Hat, you can move confidently towards SAS Viya modernization and enhance your AI and analytics journey quickly and easily.



Resources

SAS Viya

SAS Viya Install Center

Red Hat OpenShift

Red Hat OpenShift Virtualization

HPE ProLiant Gen11

HPE Alletra Storage MP B10000

HPE CSI Driver for Red Hat OpenShift

5th Gen Intel Xeon Processors

Red Hat Enterprise Linux® on Intel 5th Gen Xeon Scalable Processors

Red Hat OpenShift on Intel 4th Gen Xeon Scalable Processors

Why SAS Viya on HPE GreenLake? video

SAS Viya on HPE GreenLake: Architecture, sizing, and configuration video

SAS Viya on HPE GreenLake: Performance video

Learn more at

HPE.com/hpe-greenlake-block-storage

Chat now (sales)

© Copyright 2024 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. VMware and VMware vSphere are registered trademarks or trademarks of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. Intel and Intel Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. NVIDIA is a trademark and/or registered trademark of NVIDIA Corporation in the U.S. and other countries. Red Hat is a registered trademark of Red Hat, Inc. in the United States and other countries. Java is a registered trademark of Oracle and/or its affiliates. All third-party marks are property of their respective owners.

a00138240ENW, Rev. 3

Hewlett Packard Enterprise