



# System Requirements for SAS<sup>®</sup> Viya<sup>®</sup>

## 2020.1.5

This document might apply to additional versions of the software. Open this document in [SAS Help Center](#) and click on the version in the banner to see all available versions.

---

## Virtual Infrastructure Requirements

---

### Help with Cluster Setup

SAS Viya deployment requires experience with Kubernetes. However, SAS provides tools to help administrators create and configure a cluster that meets SAS Viya system requirements.

The SAS Viya Infrastructure as Code (IaC) projects contain scripts and configuration files that can automatically provision the infrastructure components that are required to deploy SAS Viya on Microsoft Azure, on Amazon Web Services (AWS), and on Google Cloud Platform (GCP). Each toolkit was developed to meet the same system requirements that are documented here.

---

**Note:** Some knowledge of Kubernetes and relevant third-party tools is still required. Use of the IaC tools means that some of the procedures in this guide do not apply or require modifications in order to deploy SAS Viya.

---

For more information, see these GitHub projects:

## 2

- [SAS Viya 4 Infrastructure as Code \(IaC\) for Microsoft Azure](#)
- (For the EA Program only) [SAS Viya 4 Infrastructure as Code \(IaC\) for AWS](#)
- (For the EA Program only) [SAS Viya 4 Infrastructure as Code \(IaC\) for GCP](#)

---

## Required Permissions

Deploying the software and running tasks to operate SAS Viya servers require administrative access to the cluster and to the one or more namespaces that are used. For example, any user who issues `kubectl` commands to operate or check the status of SAS Viya servers will need this level of access. In the SAS Viya documentation, this level of access is referred to as *elevated Kubernetes permissions*.

For more information, see the [Kubernetes documentation on role-based access control](#).

---

## Host Machine and Cluster Requirements

### Virtual Private Cloud Considerations

SAS Viya supports multiple VPCs and multiple subnets per deployment. Although public IP addresses are required only for the load balancer and the Ingress, SAS Viya Pods and services consume many private IP addresses. When you plan your deployment, dedicate at least one VPC with subnets that are configured with sufficiently broad CIDR ranges to accommodate these private IP addresses.

Your Kubernetes service has preferred methods for scaling up deployments with multiple IP addresses. For example, GKE supports [VPC-native clusters](#), in which Pod IP addresses are not dependent on static routes. For AKS, SAS has tested with a private CIDR and added a secondary CIDR to support larger deployments with multiple SAS Viya offerings. With EKS, consider setting up a very broad CIDR or a separate secondary IP address range in an additional subnet. For all platforms, consider using separate CIDR blocks for Pods and for services.

A CNI plug-in can be a helpful addition for IP address management in a Kubernetes cluster.

## Kubernetes Client Machine Requirements

The deployment requires a machine from which the Kubernetes command-line interface, `kubectl`, manages a Kubernetes cluster. This machine can be running Linux, Windows, or Macintosh. It requires the following client software:

- `kubectl` 1.18.8 or later

Run the following command to verify the version of `kubectl` on the client machine:

```
kubectl version --short
```

After `kubectl` is installed, note the location of the Kubernetes configuration file for use during the deployment process. The default path on Linux is `~/.kube/config`.

- `kustomize` 3.7.0

Kustomize is a client tool that is used to generate Kubernetes manifest files. Each SAS Viya release and cadence is optimized for and tested with a single version of kustomize.

Run the following command to verify the version of kustomize on the client machine:

```
kustomize version --short
```

Download kustomize here: <https://github.com/kubernetes-sigs/kustomize/releases/tag/kustomize%2Fv3.7.0>.

- (For SAS Viya Deployment Operator only) A recent version of Docker or another container runtime

The SAS Viya Deployment Operator provides a container image that includes a supported version of kubectl and kustomize. If you do not use the operator for your deployment, you can also get these versions by following the instructions in the `kubernetes-tools` README. In your `$deploy` directory, find it in `sas-bases/examples/kubernetes-tools/README.md` or in HTML format in `sas-bases/docs/using_kubernetes_tools_from_the_sas-orchestration_image.htm`.

## Node Requirements

The nodes in your Kubernetes cluster have the following requirements:

- A Linux operating system
- A 64-bit x86\_64 chipset
- A recent version of Docker or another container runtime
- Kubernetes 1.18.x or later installed on each node.

To check your version of Kubernetes, run the following command:

```
kubectl version
```

- One of the following requirements for Elasticsearch:
  - A cluster-wide setting to enable [privileged containers](#)
  - Increased virtual memory settings on the nodes that host stateful workloads

For more information about these requirements, see [“Requirements for Elasticsearch”](#) on page 6.

Additional machine requirements are described in [“Resource Guidelines”](#) on page 14.

## Kubernetes Cluster Requirements

---

**Note:** SAS Viya is supported on Amazon Web Services (AWS) and on Google Cloud Platform (GCP) only for the Early Adopter program at this time.

---

The Kubernetes cluster is where your SAS Viya software runs. The cluster has the following requirements:

- One of the following supported versions of Kubernetes:
  - (For Microsoft Azure) Microsoft Azure Kubernetes Service (AKS) 1.18.x and later
  - (For AWS) Amazon Elastic Kubernetes Service (Amazon EKS) 1.18 and later
  - (For GCP) Google Kubernetes Engine (GKE) 1.18 and later

---

**Note:** Kubernetes 1.18 reaches end-of-life status on 12 May 2021. For more information, see [Kubernetes Patch Releases](#).

---

- (For AWS) A Kubernetes Metrics Server

---

**Note:** This component is included automatically with AKS and with GKE.

---

SAS Viya was tested with version 0.4.1 of the Metrics Server. The version that is included with AKS and GKE is also supported. For more information, see <https://github.com/kubernetes-sigs/metrics-server>.

- A dedicated namespace per deployment

You can determine what namespaces are defined in your cluster by running the following command:

```
kubectl get namespaces
```

---

**Note:** Sharing a namespace with other software deployments is not generally supported. However, a namespace can be shared with software that SAS recommends in order to support logging, monitoring, and other features. Uninstalling SAS Viya causes the dedicated namespace to be deleted.

---

- An Ingress controller that is supported by Kubernetes

[Kubernetes NGINX Ingress Controller](#) version [v0.41.0 \(Helm chart version 3.8.0\)](#) or later is supported.

Istio Ingress and Istio Gateway are not supported at this time.

Some additional requirements apply. For more information, see [“Additional Configuration for Ingress and Load Balancers” on page 5](#).

- Node pools (for Microsoft Azure and GCP) or managed node groups (for AWS)

Cloud providers use different terminology to refer to collections of nodes. A minimum of three *dynamic node pools* is required:

- One default node pool

Because SAS Viya software is not intended to run on these nodes, minimal resources are required.

- Two user node pools

One of these node pools must be fully dedicated to the CAS server components. After the deployment has completed, if you find that SAS Viya software is landing on the default node pool in your configuration, reevaluate the size of the VMs in these user node pools. Sizing guidance is provided in [“Resource Guidelines” on page 14](#).

SAS recommends creating six node pools, including the default node pool.

For Microsoft Azure, node pools that span multiple availability zones are not recommended and require additional configuration. For example, Azure managed disks (the default storage class for AKS) are not zone-redundant. In this case, SAS recommends that you follow the advice in [this Microsoft document](#).

- cert-manager version 1.1.0 or later

---

**Note:** Do not install more than one instance of cert-manager per cluster.

---

You can download cert-manager from the [Jetstack GitHub site](#). Run the following command to find the currently installed version on your machine:

```
kubectl -n cert-manager describe deployments/cert-manager | grep Image
```

- **Kubernetes LoadBalancer services**

A Kubernetes LoadBalancer service is created for the Ingress Controller when it is installed. It triggers the creation of an external IP address that is associated to an external load balancer, both of which are provisioned by the cloud provider. This IP address is used for all external HTTP/HTTPS connections to SAS Viya user interfaces and services. See the [section on the LoadBalancer service](#) in the Kubernetes documentation for more information.

Additional Kubernetes load balancers are required to enable external user access to the CAS controller or to SAS/CONNECT. For more information, see [“Additional Configuration for Ingress and Load Balancers” on page 5](#).

The optional SAS Viya Infrastructure as Code (IaC) projects create the Kubernetes cluster and set up the recommended node pools. They also install the NGINX Ingress, cert-manager, and a load-balancer service. They can add these items to a cluster that they have set up, or to a cluster that you have set up as long as it complies with the requirements listed here.

## Additional Configuration for Ingress and Load Balancers

The following additional settings are required for your Ingress Controller:

- A static public IP address, which is created automatically when the Kubernetes LoadBalancer service is created for the Ingress Controller.
- The public IP address or an equivalent address for the Ingress must be registered with your DNS provider as an external endpoint.

For Microsoft Azure and for GCP, an `A` record must point to the IP address.

For AWS, a `CNAME` record must point to the fully-qualified domain name (FQDN). AWS assigns `A` records (DNS FQDNs) to load balancers. For AWS, you should create the desired DNS name for each load balancer as a `CNAME` record and configure it with the FQDN.

- An external URL that is configured for the Ingress and that is reachable from pods that run inside the cluster.

Firewall rules that allow these outbound connections are required.

- A reasonable time-out setting on the Ingress controller.

The default time-out value might be too low. SAS testing has shown successful results with a 300-second time-out in multiple environments.

- Adequate buffer sizes for request and response headers.

SAS Viya passes OAuth tokens in the request and response headers. The size of each token varies, depending on the number of group memberships that are associated with each user account. For NGINX, the default configuration uses 32 KB request headers and 16 KB response headers.

An Ingress Controller is not sufficient to enable user access to the CAS controller binary port or to the SAS/CONNECT port from outside the cluster. The connections are not HTTP connections. Therefore, you must define a Kubernetes service of type *LoadBalancer* in order to make the non-HTTP ports externally accessible.

The following settings are required for your additional Kubernetes LoadBalancer services:

## 6

- A static public IP address.
- Static DNS names and ports.
- (For AWS only) A reasonable time-out value.

The default time-out value might be too low. For example, the default time-out for AWS load balancers, 60 seconds, is too low. SAS recommends setting a 300-second time-out.

This setting can be changed for all of your CAS load balancers by specifying the `service.beta.kubernetes.io/aws-load-balancer-connection-idle-timeout` annotation in the metadata of the CASDeployment serviceTemplate. An example of the yml for this modification is included in `sas-bases/examples/cas/configure/cas-enable-external-services.yml`.

---

## Requirements for Elasticsearch

SAS Viya includes Open Distro for Elasticsearch, which is an Apache 2.0-licensed distribution of Elasticsearch with enhanced security. SAS Viya uses the Elasticsearch distributed search cluster to support search features. To enable the Elasticsearch features, your cluster must meet one of the following pre-deployment requirements:

- A cluster-wide setting must be used to enable privileged containers.
- Virtual memory settings must be increased on nodes that will host stateful workloads.

Workload classes are described in “[Plan the Workload Placement](#)” in *SAS Viya: Deployment Guide*.

When you prepare the manifest for your deployment, in order to provision the Elasticsearch pods, you must add two transformers, `internal-elasticsearch-transformer.yml` and `sysctl-transformers.yml`, to the base `kustomization.yml` file. These pods require nodes with additional virtual memory resources. To support the Elasticsearch functionality, the `sysctl-transformers.yml` transformer uses a privileged container to set the virtual memory for the `mmapfs` directory to the required level. Privileged containers are implemented as a pod security policy setting. For more information, see <https://kubernetes.io/docs/concepts/policy/pod-security-policy/>.

If privileged containers are not allowed in your environment, the Kubernetes administrator must set the virtual memory manually before performing the SAS Viya deployment. Only the nodes that are likely to run workloads in the stateful workload class are affected by this requirement.

To increase the virtual memory setting for `mmapfs` by setting it manually:

- 1 Log in to the first stateful node as root or with a sudoers account.
- 2 Set the virtual memory using the appropriate method:
  - To set the value permanently, use your preferred text editor to modify `/etc/sysctl.conf`. Update the `vm.max_map_count` setting to 262144 and save the file.
  - To set the value temporarily, run the following command:

```
sysctl -w vm.max_map_count=262144
```

- 3 Restart the node.
- 4 (Optional) Verify the modified setting:

```
sysctl vm.max_map_count
```

- 5 Repeat the previous steps on each node where stateful workloads might be hosted.

If you are using a managed Kubernetes cluster, the nodes are most likely dynamically provisioned by your cloud provider. In that situation, be aware that manual modifications do not persist after a restart of a Kubernetes node. The cluster administrator must find an alternative method to save the `vm.max_map_count` setting.

---

# Hardware and Resource Requirements

---

## Storage Requirements

Use the information in this section to estimate the sizes of storage devices for your deployment.

## Storage Overview

SAS Viya requires persistent storage volumes. Some SAS Viya components require shared storage. In general, network-based storage performs better than locally mounted storage. However, SAS recommends provisioning the CAS server and compute nodes with both high-performing local storage and shared storage.

When selecting a storage option, keep in mind that your assortment of SAS product offerings, the number of users, and the amount of data that is processed all affect sizing requirements for performance. The performance of shared storage options is highly vendor-dependent.

Instructions for modifying storage settings are provided in README files for several SAS Viya components. For example, the file titled “Configuration Settings for CAS” explains how to change the storage size for CAS PersistentVolumeClaims (PVCs), how to modify the resource allocation for ephemeral storage, and how to change the accessModes on the CAS permstore and data PVCs. The README files are located at `$deploy/sas-bases/examples/` (for Markdown format) or at `$deploy/sas-bases/docs/` (for HTML format). To modify storage for any component, first consult its README file.

## I/O Throughput and Performance Considerations

The peak I/O throughput requirements of your SAS Viya deployment might exceed the capabilities of your storage configuration. SAS generally recommends sequential I/O bandwidth of 90-120 MB per second, per physical CPU core for both persistent and ephemeral storage. For best performance, select VM instance types with the highest available I/O throughput levels. On Microsoft Azure, Premium storage is required in order to achieve these I/O throughput levels.

Although it is high-performing, NVMe offers only ReadWriteOnce (RWO) storage, which does not persist after a pod restarts. ReadWriteMany (RWX) storage is required for multiple SAS Viya components.

You should consult your cloud provider’s storage and compute instance documentation to ensure that the storage environment can provide the level of I/O that is required. SAS Technical Support found that many performance issues reported by SAS Viya customers can be directly attributed to insufficient levels of I/O throughput.

In general, SSDs perform better than HDDs.

## Disk Space

Most SAS Viya software uses the same base container and Docker layers in order to optimize disk usage. The minimum amount of disk space that is required for the installation and for logging is 48 GB. Therefore, the minimum combined capacity of the Kubernetes worker nodes should be 48 GB.

The number of containers and the size of the images depend on the products in the software order. These images require approximately 30-60 GB of disk space on the host on which you are downloading SAS images and also in the destination registry.

If you are using the optional SAS Mirror Manager, the same sizing guidelines apply to the machine where you run it. After the first run, SAS Mirror Manager creates a local copy of the container images to make the process of mirroring images as quick as possible.

The topic titled [Virtual Machine Recommendations for Microsoft Azure](#) provides guidelines for VM sizing, including recommended disk size.

## Persistent Storage Volumes, PersistentVolumeClaims, and Storage Classes

Persistent storage is required by multiple SAS Viya components. Verify that at least one ReadWriteMany (RWX) StorageClass has been defined and set as your default. Run the following command to verify that a default StorageClass has been defined:

```
kubectl get storageclass
```

Make a note of the names of all storage classes in case you later need to modify storage settings. You must also update the base kustomization.yaml to specify the PVCs that are associated with the RWX StorageClass. For more information, see [Specify PersistentVolumeClaims to Use ReadWriteMany StorageClass](#).

Many storage classes have a ReclaimPolicy that is set to **Delete** by default. If you delete a namespace that includes such storage classes, the PVCs in that namespace are deleted. If the ReclaimPolicy is Delete, the corresponding persistent volumes (PVs) are also deleted, resulting in data loss. For information about changing the ReclaimPolicy in Kubernetes, see [Change the Reclaim Policy of a PersistentVolume](#).

Multiple PVCs are configured automatically during the deployment. Here are the default settings for all deployments. Note that the default settings correspond to the minimum recommended sizes:

- **cacheserver pod** — SAS Cache Server stores longer-term data, such as configuration settings, and serves it to SAS Viya services. It requires a storage volume. It is deployed with high availability (HA) by default, and two replicas, each with a 2 Gi PVC; accessMode: ReadWriteOnce (RWO).
- **Consul** — Supports SAS Configuration Server. It is deployed with high availability (HA) by default, and three replicas, each with a 1 Gi PVC; accessMode: RWO.

The Consul pod also requires a mount for the Consul data directory. The mount is an empty directory that points to `/consul/data`. It is configured automatically by the deployment process.

- **RabbitMQ** — Supports SAS Message Broker. It is deployed with HA by default and three replicas, each with a 2 Gi PVC; accessMode: RWO.
- **PostgreSQL** — Supports SAS Infrastructure Data Server. The requirements depend on whether you deploy the default (internal) PostgreSQL instance or you supply your own PostgreSQL server (external). Requirements are summarized in the following table:

**Table 1** SAS Infrastructure Data Server Storage Requirements

PostgreSQL Deployment	accessMode	Default Size
Internal (the default deployment)	RWO	Deploys with HA by default: <ul style="list-style-type: none"> <li>■ Three nodes <ul style="list-style-type: none"> <li>One primary and two replicant data nodes.</li> </ul> </li> <li>■ PVC of 128 GB per node</li> </ul>
External (you are responsible for this server)	Not applicable	A minimum of one volume, 128 GB of space.

For more information, see “[PostgreSQL Deployment Options](#)”.

- CAS server — Requirements depend on whether you deploy the CAS server as SMP or MPP. Two PVCs are required in either case:
  - cas-default-permstore — Required in order to store caslib management privileges
  - cas-default-data — The default location for new caslibs; equivalent to the SAS Viya CASDATADIR setting

Their requirements are summarized in the following table:

**Table 2** CAS Server Persistent Storage Requirements

CAS Server	accessMode	Default Size
SMP CAS server	Either RWO or RWX. Set to RWX by default.	cas-default-permstore: 100 Mi cas-default-data: 8 Gi
MPP CAS server	RWX RWX	cas-default-permstore: 100 Mi cas-default-data: 8 Gi

RWX accessMode is required for any backup controllers for CAS. To change the accessMode for either cas-default-data or cas-default-permstore, perform the steps that are described in [Change accessMode](#).

Every CAS pod also requires a mount for the CAS disk cache. The mount, which is used as the CAS disk caching space, is an empty directory that points to `/cas/cache`. It is configured automatically by the deployment process. Performance is improved when the `CASENV_CAS_DISK_CACHE` environment variable is configured to point to a striped array (RAID 0) across all internal (local or ephemeral) disks on each VM instance where a CAS pod runs. For more information, see [Tune CAS\\_DISK\\_CACHE](#).

You can modify most default resources that are provided for your CAS server. However, the mountPath that is defined for cas-default-permstore and cas-default-data cannot be modified. The associated mount points are `/cas/permstore` and `/cas/data`, respectively. To define mounts that are added to CAS overlays, you can add patchTransformers to your manifests.

- Elasticsearch — Supports the search feature. Creates one PVC. Default size: 128 Gi; accessMode: RWO.
- Backup and restore operations — Require two PVCs, described as follows:
  - sas-common-backup-data — Stores a backup file of settings for SAS Infrastructure Data Server and SAS Configuration Server. Default size: 25 GB; accessMode: RWX.
  - sas-cas-backup-data — Stores a backup of CAS server data and the CAS permstore. Default size: 8 GB; accessMode: RWX.

**IMPORTANT** These default settings are the minimum recommended sizes. Verify that the ReclaimPolicy for these PVCs is set to **Retain**.

## Persistent Volumes for Applications

Some individual SAS Viya products also require persistent storage. If your software order included these applications, you must set up additional volumes with the required settings:

- SAS Data Quality — Requires a volume for SAS Quality Knowledge Base data. Default size: 8 Gi; accessMode: RWX.
- SAS Micro Analytic Service — Requires storage volumes if the optional features, ASTORES or archives, are configured:
  - Volume for ASTORES — Default size: 30 GB; accessMode: RWX. Depending on model complexity and the number of models, more space might be required.

To resize the PVC for ASTORES, update the overlay with the new size and run it. If models have already been loaded, SAS recommends setting the storage class attribute `allowVolumeExpansion` to true. If the provider does not support the storage class, resize the volume and then republish all ASTORE models.

- Volume for archive logs — If the archive feature is enabled in SAS Micro Analytic Service, it stores input and output transaction logs in a persistent volume. Default size: 30 GB; accessMode: RWX.

This volume grows continuously. SAS recommends that you create regular backups of this data and monitor the available space.

- SAS Studio — Requires persistent storage if users will take advantage of Git integration features.

To enable Git integration, SAS Studio users require access to a shared file system that is accessible to the pod where the compute server is running (the pod where the compute workload has been placed). Be aware that storage on this pod is temporary unless you configure persistent storage. Configure persistent storage for the compute server if you plan to use the Git integration features. For more information, see [“Creating Persistent File Storage” in SAS Studio: Administrator’s Guide](#).

Work that is associated with other SAS Studio tasks is stored in SAS Content, which does not require a persistent volume.

## Additional PVC Requirements for Microsoft Azure

In a Microsoft Azure environment, any RWO persistent storage should be configured on VM instance types whose names include the “S” suffix. This suffix indicates that the VM instances have support for Premium (SSD) Storage.

By default, PVCs are owned by the root user. In Microsoft Azure environments, you must define a StorageClass in which you explicitly set mountOptions for each PVC in order to allow non-root users to access them. If you do not allow access by the sas user account, at a minimum, you will see permission errors. For more information about the sas user account, see [User Accounts on page 37](#).

In your custom StorageClass definition, the UID and GID mount options must match the container process ID (which is 1001, sas, by default). Here is an example configuration:

```
cat << 'EOF' > StorageClass-RWX.yaml
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: sas-azurefile
provisioner: kubernetes.io/azure-file
mountOptions:
  - dir_mode=0777
  - file_mode=0777
  - uid=1001
  - gid=1001
parameters:
  skuName: Standard_LRS
allowVolumeExpansion: true
EOF
kubectl apply -f StorageClass-RWX.yaml
```

Although Azure Kubernetes Service (AKS) provides a default azurefile storage class, the SAS testing organization found that this default storage class could not be modified to include the mount point options that SAS Viya requires. Instead, AKS quickly reverts the changes back to the default. Therefore, SAS recommends that you create a custom storage class in Azure environments.

## File System and Shared Storage Requirements

The persistent volumes described previously are required for component and application data. However, various SAS Viya components also require a shared file system. A shared file system is required for multiple purposes, which include shared data storage and private user directories. A file server that uses the network file system (NFS) protocol is the minimum requirement.

As you select a file system option, consider both disk size and the number of disks. Speed is directly related to the amount of disk space that is provisioned. It is also constrained by VM network and I/O limits.

In a Microsoft Azure environment, each VM instance type has a maximum input/output operations per second (IOPS) metric and a throughput metric, as well as a maximum number of disks that can be attached. The Microsoft Azure VM instances that SAS recommends are optimized for IOPS rather than for storage throughput. Striping can be used to increase storage IOPS and throughput, both of which are important for SAS Viya performance. The performance of attached disk configurations is constrained by VM limits.

The following options for shared file storage are recommended:

Table 3 Storage Recommendations

Cloud Provider	Standard Deployments	HA Deployments	Notes
Microsoft Azure	<p><a href="#">Azure Premium SSD Managed Disks</a> with a shared file storage layer.</p> <p>When volumes are mounted to a VM instance that exports them as NFS, they have the equivalent of RWX accessMode.</p> <p>RAID 5 is recommended.</p>	<a href="#">Azure NetApp Files</a>	Both of these storage options are encrypted by default.
AWS	<p>Elastic Block Store (EBS), mounted to a VM that exports with NFS.</p> <p>When volumes are mounted to a VM instance that exports them as NFS, they have the equivalent of RWX accessMode.</p>	<p>Amazon Elastic File Share (EFS), with the following options:</p> <ul style="list-style-type: none"> <li>■ Performance mode set to <b>Max I/O</b></li> <li>■ Throughput mode provisioned with <b>1024 MiB/s</b> or more</li> </ul>	<p>Installing a provisioner for EFS, such as the <code>nfs-provisioner</code> or <code>nfs-subdir-external-provisioner</code>, is recommended.</p> <p>RAID 0 is recommended.</p>
GCP	<p>Cloud Storage</p> <p>When volumes are mounted to a VM instance that exports them as NFS, they have the equivalent of RWX accessMode.</p>	Cloud Filestore, SSD tier	RAID 0 is recommended.

**Note:** SAS Viya is supported on Amazon Web Services (AWS) and on Google Cloud Platform (GCP) only for the Early Adopter program at this time.

SAS also recommends using a consistent directory structure in your shared storage solution. Consistency is helpful for the following reasons:

- Multiple SAS Viya components require data storage.
- Some SAS offerings can also use a shared location for private user directories.
- Shared storage can serve as the location for persistent volumes to be provisioned in a consistent manner.

- SAS Viya uses some shared, open-source binary files.

To optimize your deployment, create the NFS directory structure that is described in the following table:

**Table 4** Shared File System Recommended Directory Structure

Directory	Description
<code>/astores</code>	Location of the shared directory for ASTORES and ASTORE models.
<code>/bin</code>	Location for open-source companion software directories.
<code>/bin/nfsviapython</code>	Location for your Python binary files.
<code>/bin /nfsviya-r</code>	Location for your R binary files.
<code>/data</code>	Location for SAS and CAS data.
<code>/homes</code>	Location for user private directories. As a best practice, create a subdirectory for each user of SAS Viya offerings.
<code>/pvs</code>	Location for persistent volumes that are provisioned using the file system.
<code>/permstore</code>	Location for the CAS server to store caslib authorization information.
<code>/backups</code>	Locations where SAS Viya backup files can be saved.
<code>/backups/cas</code>	
<code>/backups/common</code>	
<code>/quality-knowledge-base</code>	Location where SAS Quality Knowledge Base stores data.

## Requirements for GPU Support

Deep learning capabilities are included with SAS Visual Machine Learning. The deep learning features are automatically installed. A graphical processing unit (GPU) is not required in your cluster in order to use deep learning features. However, a GPU provides additional functionality.

To enable deep learning with GPU functionality, here are additional requirements on Linux operating systems:

- A powerful, CUDA-capable GPU.

SAS has tested with GPUs that have the NVIDIA Compute Unified Device Architecture (CUDA).

Only the NVIDIA Quadro and Tesla product families are supported.

NVIDIA Pascal, Volta, Turing, and Ampere architectures are supported.

---

**Note:** If you have multiple GPUs installed on the same node, they must be of the same model.

---

- The NVIDIA display driver, version 440.33.01 or later. SAS recommends using the latest version.

SAS also recommends enabling NVIDIA driver persistence mode at all times. For more information, see the [Driver Persistence](#) section of the NVIDIA deployment documentation.

When you install the NVIDIA display driver on Linux, SAS recommends following the instructions in the NVIDIA CUDA [Installation Guide for Linux](#).

You can download the current drivers from <http://www.nvidia.com/Download/index.aspx?lang=en-us>.

- `/lib64` is the first path that is defined for the `LD_LIBRARY_PATH` environment variable on the server where the GPU is installed.

Run the following command on the machine where the GPU is installed in order to check the device type, the driver version, and the CUDA version:

```
nvidia-smi
```

---

## Resource Guidelines

The topics in this section provide recommendations for node scheduling, node sizing, and performance optimization.

## Factors That Affect Resource Requirements

Several factors affect resource utilization by SAS Viya components, such as the following:

- Whether your CAS server implementation is SMP or MPP  
Your deployment of SAS Cloud Analytic Services (CAS) can be performed on a single node (SMP) or across several nodes (MPP). Distributing the CAS server across multiple nodes enables massively parallel processing (MPP). MPP CAS typically requires more resources.
- The expected amount of data that SAS Viya users will process
- The expected number of concurrent users
- Whether an optional GPU is used in order to leverage the SAS Deep Learning feature

## Sizing for Migration

If you are migrating to this version of SAS Viya from a previous (3.x) version, the resource utilization of both deployments will be similar.

Here is an example: if your SAS Viya 3.5 deployment included three CAS server machines with 256 GB of RAM each, you should reserve nodes with a total of 768 GB of RAM in the Kubernetes cluster. If the Kubernetes administrator has allocated physical nodes of 64 GB of RAM each, you should label 12 of those nodes for CAS servers.

However, accounting for the typical 90% resource utilization, you can assume that the CAS controller and workers can reliably access only about 691 GB of RAM in this environment. Therefore, one additional node with 64 GB of RAM is recommended.

## Sizing Recommendations for Microsoft Azure

A minimum of three node pools is required in your cluster: one default (system) node pool, where Kubernetes and other components that are not SAS Viya are deployed, and two user node pools for SAS Viya. Refer to “[Default Node Pool Configuration](#)” in *SAS Viya: Deployment Guide* for instructions on configuring the system node pool. One user node pool must be fully dedicated to the CAS server. The nodes in the CAS node pool require [taints](#) to prevent Kubernetes from scheduling non-CAS workloads on them. The other user node pool can host the remaining (non-CAS) component workloads. Consider creating five user node pools in order to accommodate the five workload classes.

The following table provides baseline resource recommendations for representative SAS Viya product offerings. Derived from SAS performance testing, these estimates are for a *medium-sized* deployment, which is defined as 10 concurrent users that access SAS Viya user interfaces. In SAS testing simulations, these users manipulated data sets of 10–30 GB. VM instances all used an Intel Cascade Lake processor. The instance types are hyper-threaded, which means that two vCPUs are equivalent to one physical CPU core.

**Table 5** Baseline Resource Requirements per Offering

Offering	CAS Node Pool	System Node Pool	User Node Pools*
SAS Visual Analytics and SAS Data Preparation	RAM: 128 GB per instance CPU: 16 vCPUs per instance Example: Microsoft Azure E16ds_v4 Recommended minimum for CAS disk cache: 150 GB; ephemeral storage Example number of machines: 4	RAM: 64 GB CPU: 8 vCPUs Example: Microsoft Azure Standard_E8ds_v4 Example number of machines: 1	RAM: 64 GB CPU: 8 vCPUs per instance Example: Microsoft Azure E16ds_v4 Recommended minimum ephemeral storage for selected applications: 60 GB Example number of machines: 1 per node pool
SAS Visual Machine Learning	RAM: 128 GB per instance CPU: 16 vCPUs per instance Example: Microsoft Azure E16ds_v4 Recommended minimum for CAS disk cache: 600 GB; ephemeral storage Example number of machines: 8	RAM: 64 GB CPU: 8 vCPUs Example: Microsoft Azure Standard_E8ds_v4 Example number of machines: 1	RAM: 128 GB per instance CPU: 16 vCPUs per instance Example: Microsoft Azure E16ds_v4 Minimum disk: 2 x 128 GB Example number of machines: 1 per node pool

Offering	CAS Node Pool	System Node Pool	User Node Pools*
SAS Visual Data Science	RAM: 384 GB per instance CPU: 48 vCPUs per instance Example: Microsoft Azure E48ds_v4 Recommended minimum for CAS disk cache: 1200 GB Example number of machines: 9	RAM: 64 GB CPU: 8 vCPUs per instance Example: Microsoft Azure Standard_E8ds_v4 Example number of machines: 1	RAM: 128 GB per instance CPU: 48 vCPUs per instance Example: Microsoft Azure E16ds_v4 Minimum disk: 2 x 128 GB Example number of machines: 1 per node pool

\* In addition to the CAS node pool, four user node pools are recommended in order to host the remaining SAS Viya workload classes.

**Note:** The additional resources that are recommended for SAS Visual Machine Learning are needed for compute workloads. The compute class includes pods that host SAS compute server instances, or that run batch jobs, or that do both.

The SAS workload node placement strategy suggests that you dedicate a node or a node pool on which Kubernetes schedules, or prefers scheduling, only SAS compute workloads. These workloads consist of pods that host SAS compute server instances, that run batch jobs, or that do both. This strategy enables more granular tuning when you want a specialized set of nodes to manage the compute workloads in the environment. SAS recommends selecting a machine with additional resources for the compute workload class. For more information about the workload node placement strategy, see [“Plan the Workload Placement” in SAS Viya: Deployment Guide](#).

By default, updates are applied to your deployment using a strategy that avoids downtime. During each SAS Viya software update, the existing pods continue running until the new, updated pods have started up. Therefore, your cluster requires sufficient available resources to enable duplicates of each pod to run temporarily.

These guidelines do not attempt to account for all ordering scenarios, but instead are intended to illustrate typical software orders. SAS strongly recommends that you consult with a sizing expert to obtain an official hardware recommendation that is based on your requirements. To request sizing expertise, contact your SAS account representative. If you need assistance in determining your SAS account representative, send an email to [contactcenter@sas.com](mailto:contactcenter@sas.com).

## Sizing Recommendations for AWS

**Note:** SAS Viya is supported on AWS only for the Early Adopter program at this time.

A minimum of three AWS managed node groups is required in your cluster. One node group should be reserved for components that are not SAS Viya, and the remaining managed node groups are required for SAS Viya components. The reserved node group is referred to as the “default node pool” in this document. Refer to [“Default Node Pool Configuration” in SAS Viya: Deployment Guide](#) for instructions on configuring it. One managed node group must be fully dedicated to the CAS server. The nodes in the CAS node group require [taints](#) to prevent Kubernetes from scheduling non-CAS workloads on them. The other managed node group can be dedicated to the remaining (non-CAS)

component workloads. Consider creating six managed node groups in order to accommodate the five workload classes and the default node pool.

The following table provides baseline resource recommendations for a representative SAS Viya product offering. Derived from SAS performance testing, these estimates are for a *medium-sized* deployment, which is defined as 10 concurrent users that access SAS Viya user interfaces. In SAS testing simulations, these users manipulated data sets of 10–30 GB. The EC2 VM instances used an Intel Cascade Lake series processor. In AWS, the number of vCPUs for an instance is indicated by a number in the name of the instance type, multiplied by 4. For example, the m5n.2xlarge instance type provides 8 vCPUs, the equivalent of 4 physical CPU cores.

**Table 6** *Baseline Resource Requirements for an Example Offering*

Workload Class	SAS Visual Machine Learning
Default Node Pool	RAM: 16 GB CPU: 4 vCPUs Example: m5n.xlarge Example number of machines: 1
CAS	RAM: 128 GB per instance CPU: 48 vCPUs per instance Example: r5dn.12xlarge Storage for CAS disk cache: 2 x 900 NVMe SSD Example number of machines: 3 (MPP) or 1 (SMP)
Connect	RAM: 128 GB per instance CPU: 16 vCPUs Example: r5dn.4xlarge Example number of machines: 2
Compute	RAM: 384 GB per instance CPU: 48 vCPUs per instance Example: r5dn.12xlarge Storage for SASWork directory: 2 x 900 NVMe SSD Example number of machines: 2
Stateful or Stateless	RAM: 64 GB per instance CPU: 16 vCPUs per instance Example: m5n.4xlarge Example number of machines: 2 per node group

The SAS [workload node placement strategy](#) suggests that you dedicate a node on which Kubernetes schedules, or prefers scheduling, only SAS compute workloads. These workloads consist of pods that host SAS compute server instances, that run batch jobs, or that do both. This strategy enables more granular tuning when you want a specialized set of nodes to manage the compute workloads in the environment.

By default, updates are applied to your deployment using a strategy that avoids downtime. During each SAS Viya software update, the existing pods continue running until the new, updated pods have started up. Therefore, your cluster requires sufficient available resources to enable duplicates of each pod to run temporarily.

These guidelines do not attempt to account for all ordering scenarios, but instead are intended to illustrate typical software orders. SAS strongly recommends that you consult with a sizing expert to obtain an official hardware recommendation that is based on your requirements. To request sizing expertise, contact your SAS account representative. If you need assistance in determining your SAS account representative, send an email to [contactcenter@sas.com](mailto:contactcenter@sas.com).

## Sizing Recommendations for GCP

---

**Note:** SAS Viya is supported on GCP only for the Early Adopter program at this time.

---

A minimum of three GCP node pools is required in your cluster. One node pool should be reserved for components that are not SAS Viya, and the remaining node pools are required for SAS Viya components. The reserved node pool is referred to as the “default node pool” in this document. Refer to [“Default Node Pool Configuration” in SAS Viya: Deployment Guide](#) for instructions on configuring it. One node pool must be fully dedicated to the CAS server. The nodes in the CAS node group require [taints](#) to prevent Kubernetes from scheduling non-CAS workloads on them. The other node pool can be dedicated to the remaining (non-CAS) component workloads. Consider creating six node pools in order to accommodate the five workload classes and the default node pool.

The following table provides baseline resource recommendations for a representative SAS Viya product offering. Derived from SAS performance testing, these estimates are for a *medium-sized* deployment, which is defined as 3 concurrent users that access SAS Viya user interfaces. In SAS testing simulations, these users worked with table sizes of 76 GB and row counts of 10,000,000. The VM instances used an Intel Cascade Lake series processor.

**Table 7** Baseline Resource Requirements for an Example Offering

Workload Class	SAS Visual Machine Learning
Default Node Pool	RAM: 32 GB CPU: 8 vCPUs Example: n2-standard-8 Example number of machines: 1
CAS	RAM: 384 GB CPU: 48 vCPUs Example: n2-highmem-48 Storage for CAS disk cache: 3 TiB; ephemeral Storage for CAS default data: 2.5 TiB, basic SDD (RAID 0) Example number of machines: 3 (MPP) or 1 (SMP)
Connect	RAM: 128 GB per instance

Workload Class	SAS Visual Machine Learning
	CPU: 16 vCPUs Example: n2-highmem-16 Example number of machines: 1
Compute	RAM: 384 GB per instance CPU: 48 vCPUs per instance Example: n2-highmem-48 Storage for SASWork directory: 5 x 375 GB local SSD (RAID 0) Example number of machines: 1
Stateful or Stateless	RAM: 64 GB per instance CPU: 16 VCPUs per instance Example: n2-standard-16 Recommended minimum ephemeral storage for selected applications: 60 GB Example number of machines: 2

The SAS [workload node placement strategy](#) suggests that you dedicate a node on which Kubernetes schedules, or prefers scheduling, only SAS compute workloads. These workloads consist of pods that host SAS compute server instances, that run batch jobs, or that do both. This strategy enables more granular tuning when you want a specialized set of nodes to manage the compute workloads in the environment.

By default, updates are applied to your deployment using a strategy that avoids downtime. During each SAS Viya software update, the existing pods continue running until the new, updated pods have started up. Therefore, your cluster requires sufficient available resources to enable duplicates of each pod to run temporarily.

These guidelines do not attempt to account for all ordering scenarios, but instead are intended to illustrate typical software orders. SAS strongly recommends that you consult with a sizing expert to obtain an official hardware recommendation that is based on your requirements. To request sizing expertise, contact your SAS account representative. If you need assistance in determining your SAS account representative, send an email to [contactcenter@sas.com](mailto:contactcenter@sas.com).

## CAS Server Resources

The CAS server consists of at least one CAS controller and at least one CAS worker, which can share a single node (SMP) or can be deployed on multiple nodes (MPP or distributed). The nodes that you select for the CAS controller and for each worker have the following requirements:

- Guaranteed [quality of service](#) (QoS)

QoS is required in order to avoid pod shutdowns when memory resources are limited. Kubernetes QoS requires that the requests and limits be set to equal numbers for each container in the pod.

CAS pods must execute with guaranteed QoS. The CAS controller ensures guaranteed QoS when the proper customizations are applied (specifically, `sas-bases/overlays/cas-server/auto-resources` and `sas-bases/overlays/cas-server/auto-resources/remove-`

`resources.yaml`). The file `sas-bases/examples/cas/configure/cas-manage-cpu-and-memory.yaml` can be used as a template to manually specify resource requests and limits.

- A mount for the CAS disk cache

Required mounts and PVCs are described in “[Persistent Storage Volumes, PersistentVolumeClaims, and Storage Classes](#)” on page 8. Performance improves if the CAS server runs on a node that includes a mounted SSD storage volume.

- Fully dedicated nodes that do not share resources with any other SAS Viya components

To create these dedicated resources, apply taints and labels to the appropriate nodes in the CAS node pool. This recommendation ensures that CAS QoS can be achieved. The documented node taints prevent Kubernetes from scheduling workloads that could compete for resources with CAS. See [Plan Workload Placement](#) for more information.

If resources on CAS-dedicated nodes are exhausted, CAS might be scheduled to nodes outside of the CAS node pool. If external scheduling occurs, contact your SAS account representative to address issues with CAS node pool sizing. To avoid this situation, SAS recommends that you add taints to other nodes in the cluster in conformance with the SAS workload placement strategy.

**IMPORTANT** Tainting all nodes in the cluster limits workloads from components that are not SAS components from being scheduled to tainted nodes. Otherwise, adjustments to manifests are necessary for workloads that are not SAS components.

## Virtual Machine Recommendations for Microsoft Azure

The following table provides guidelines to help you set up an appropriate cloud environment for your Kubernetes deployment of SAS Viya. Keep in mind that these are minimum recommended allocations. Monitoring workload levels and available storage is recommended in order to maintain adequate resources for the cluster.

*Table 8 Cloud Components and Minimum Sizing*

Component	Description	Size
Nodes	<p>Select virtual machines with good I/O, such as the D(s/ds)v3 Series machines. These instances provide hyper-threading and up to 3.5 GHz.</p> <p>Use the following general guidance:</p> <ul style="list-style-type: none"> <li>■ Select the E16ds_v4 or equivalent for the CAS server.</li> <li>■ Select Standard_D16s_v3 for most SAS Viya orders.</li> <li>■ Select Standard_D8s_v3 for SAS Viya orders with fewer offerings (1–2).</li> </ul>	<p>One large deployment: 5 or more nodes</p> <p>Two or more large deployments: 7 or more nodes</p> <p>With autoscaling, the following are recommended:</p> <ul style="list-style-type: none"> <li>■ 2 or more nodes in each of the user node pools</li> <li>■ 1 node in the system node pool</li> </ul>

Component	Description	Size
	<ul style="list-style-type: none"> <li>Select the auto-scaling option for node pools where possible.</li> </ul>	
Disk Size	In Microsoft Azure, each node includes ephemeral storage.	200 GB is the recommended minimum. The default 100 GB disk size in Azure was not adequate in SAS testing.
Database Connections	A PostgreSQL database is required to host SAS Infrastructure Data Server. You can use the default <i>internal</i> database that is deployed automatically, or you can supply your own <i>external</i> database. For more information, see <a href="#">“Internal versus External PostgreSQL Instances” on page 39</a> .	A typical deployment with two or three SAS Viya offerings requires a minimum of 1024 database user connections to avoid running out of connections. The internal PostgreSQL instance is set to 1280 connections by default. An external PostgreSQL server should support maximum connections and maximum prepared transactions of at least 1024.
Database Storage	A single PostgreSQL server can support multiple SAS Viya deployments. You define a separate database for each deployment.	<p>A typical deployment with two or three SAS Viya offerings requires approximately 1–2 GB for the deployment process.</p> <p>If a single PostgreSQL server supports multiple SAS Viya deployments that are typical in size, use the following guidance to determine storage space:</p> <ul style="list-style-type: none"> <li>For 1–2 deployments, use at least 128 GB of storage space.</li> <li>For 3–4 deployments, use at least 512 GB of storage space.</li> </ul>

Consider the following guidelines as you select an image type for the VM instances in your cluster:

- Select VM instance types with Intel chips. If the name of the Azure image contains the letter “a,” the VM is based on an AMD chip (such as the Dav4 and Dasv4-series). Multiple SAS Viya processes performed significantly faster on Intel chips in SAS testing.
- For the v4 type of image, if the name contains the letter “d” indicating “disk,” the node includes a local temp disk. The CAS server requires ephemeral storage.
- If the node is intended to host a pod that requires a PVC, SAS recommends that you select a VM image that supports premium storage. Premium storage is indicated by the letter “s” in the image name. For more information, see [“Storage Requirements” on page 7](#).

- SAS Viya performs better on nodes with a high memory-to-CPU ratio. The Microsoft Azure E series images have a higher memory-to-CPU ratio than the D series.

## Virtual Machine Recommendations for AWS

**Note:** SAS Viya is supported on AWS only for the Early Adopter program at this time.

The following table provides guidelines to help you set up an appropriate cloud environment for your EKS deployment of SAS Viya. Keep in mind that these are minimum recommended allocations. Monitoring workload levels and available storage is recommended in order to maintain adequate resources for the cluster.

**Table 9** Cloud Components and Minimum Sizing

Component	Description	Size
Nodes	<p>Select virtual machines with good I/O, such as the r5d or i3 series machines. These instances provide hyper-threading and up to 3.5 GHz.</p> <p>Use the following general guidance:</p> <ul style="list-style-type: none"> <li>■ Select an r5d or i3.8xlarge equivalent for the CAS server.</li> <li>■ For a medium-sized SAS Viya order with multiple offerings, select r5d.8xlarge for the remaining nodes.</li> <li>■ Select the auto-scaling option for node groups where possible.</li> </ul>	<p>One large deployment: 5 or more nodes</p> <p>Two or more large deployments: 7 or more nodes</p> <p>With autoscaling, the following are recommended:</p> <ul style="list-style-type: none"> <li>■ 2 or more nodes in each of the SAS Viya node groups</li> <li>■ 1 node in the default node group</li> </ul>
Disk Size	In AWS, each node includes ephemeral storage.	200 GB is the recommended minimum. The default instance store size in AWS is typically adequate.
Database Connections	A PostgreSQL database is required to host SAS Infrastructure Data Server. You can use the default <i>internal</i> database that is deployed automatically, or you can supply your own <i>external</i> database. For more information, see <a href="#">“Internal versus External PostgreSQL Instances” on page 39</a> .	A typical deployment with two or three SAS Viya offerings requires a minimum of 1024 database user connections to avoid running out of connections. The internal PostgreSQL instance is set to 1280 connections by default. An external PostgreSQL server should support maximum connections and maximum

Component	Description	Size
		prepared transactions of at least 1024.
Database Storage	A single PostgreSQL server can support multiple SAS Viya deployments. You define a separate database for each deployment.	<p>A typical deployment with two or three SAS Viya products requires approximately 1–2 GB for the deployment process.</p> <p>If a single PostgreSQL server supports multiple SAS Viya deployments that are typical in size, use the following guidance to determine storage space:</p> <ul style="list-style-type: none"> <li>■ For 1–2 deployments, use at least 128 GB of storage space.</li> <li>■ For 3–4 deployments, use at least 512 GB of storage space.</li> </ul>

Consider the following guidelines as you select instance types for the VM instances in your cluster:

- Select VM instance types with Intel chips. Multiple SAS Viya processes performed significantly faster on Intel chips in SAS testing.
- The CAS server requires ephemeral storage. If the name of the AWS instance type contains the letter “d” indicating “disk,” the node includes a local temp disk.
- SAS Viya performs better on nodes with a high memory-to-CPU ratio.

## Data Source Requirements

### Data Source Requirements

Your software order included SAS/ACCESS products to support your data sources. Data access requires additional customization of your Kubernetes deployment and in some cases also requires installation of other software.

Individual SAS Viya product offerings might support a subset of these data sources. Refer to the lists of data sources that are supported by individual offerings. Then consult the list of requirements for the SAS/ACCESS product that supports your data source for additional system requirements that apply to your environment.

## Supported Data Sources

The following external data sources are supported by most SAS Viya offerings, including SAS Visual Analytics, SAS Visual Statistics, SAS Visual Machine Learning, SAS Visual Data Science, SAS Visual Data Science Decisioning, SAS Visual Forecasting, and SAS Visual Text Analytics:

- Amazon Redshift
- Google BigQuery
- Greenplum
- Hadoop with Hive
- IBM DB2
- IBM Netezza
- Impala
- Data sources accessible with a Java Database Connectivity (JDBC) driver
- Microsoft SQL Server
- MongoDB
- MySQL
- Data sources accessible with an ODBC driver
- Oracle
- PC files
- PI System
- PostgreSQL
- Salesforce
- SAP ASE
- SAP HANA
- SAP R/3
- Snowflake
- Spark
- Teradata
- Vertica
- Yellowbrick

## Data Sources for SAS Intelligent Decisioning and SAS Micro Analytic Service

SAS Intelligent Decisioning and SAS Micro Analytic Service support only the following subset of the data sources that SAS Viya supports:

- Microsoft Azure PostgreSQL

- Microsoft SQL Server
- Oracle
- PostgreSQL

---

## General Requirements for SAS/ACCESS

Before you start the deployment, collect the third-party libraries and configuration files that are required for your data sources. Examples of these requirements include the following:

- Third-party drivers
- ODBC drivers
- JDBC drivers
- Hadoop configuration files

When you have collected these files, place them on storage that is accessible to your Kubernetes cluster. This storage could be a mount or a storage device with a persistent volume configured.

SAS recommends organizing your software in a consistent manner on your mount or storage device. Take note of the details for your specific storage solution, as well as the paths to the configuration files within it. You will need this information before you start the deployment.

---

## Requirements for SAS/ACCESS Interface to Amazon Redshift

SAS/ACCESS Interface to Amazon Redshift includes SAS Data Connector to Amazon Redshift.

The required client software is included with your SAS Viya installation. Using a Data Source Name (DSN) to connect to Amazon Redshift requires post-installation configuration of your Kubernetes deployment.

---

## Requirements for SAS/ACCESS Interface to DB2

SAS/ACCESS Interface to DB2 includes SAS Data Connector to DB2.

IBM DB2 Connect™ must also be licensed if you plan to connect to IBM DB2 databases that are running on AS/400, VSE, VM, MVS, and z/OS systems. The following DBMS products are supported:

- IBM DB2 version 11.5 or later
- IBM Integrated Analytics System (IIAS)
- Client utilities for IBM DB2 version 11.5 or later

SAS recommends installing the latest FixPack on the client and server.

---

## Requirements for SAS/ACCESS Interface to Google BigQuery

SAS/ACCESS Interface to Google BigQuery includes SAS Data Connector to Google BigQuery.

Required client software is included with your SAS Viya installation. No additional software is required.

---

## Requirements for SAS/ACCESS Interface to Greenplum

SAS/ACCESS Interface to Greenplum includes SAS Data Connector to Greenplum.

SAS/ACCESS Interface to Greenplum supports Greenplum Database versions 6.5 or later, and the required client software is included with your SAS Viya installation. Using a Data Source Name (DSN) to connect to Greenplum requires post-installation configuration of your Kubernetes deployment.

---

## Requirements for SAS/ACCESS Interface to Hadoop

SAS/ACCESS Interface to Hadoop includes SAS Data Connector for Hive.

Hive 1.1 or later is required.

SAS/ACCESS Interface to Hadoop supports the following Hadoop distributions:

- Cloudera CDH 6.2
- Cloudera Data Platform (CDP) 7.1 (Public and Private Cloud)
- Hortonworks HDP 3.1
- Amazon Web Services EMR 5.13
- Microsoft Azure HDInsight 3.6 and 4

The SAS policy that applies to alternative releases or distributions of Hadoop is documented [on the SAS Support website](#).

You must run the Hadoop Tracer Script for SAS Viya in order to install required JAR files that enable the CAS server to access files in Hadoop. The script has the following requirements:

- The user who is running the script must have authorization to issue HDFS and Hive commands.
- Python and the strace Linux library must be installed on the Hadoop cluster. Install them from the package repositories for your Linux distribution if necessary.

Python 2.6 or later is required.

---

## Requirements for SAS/ACCESS Interface to Impala

SAS/ACCESS Interface to Impala includes SAS Data Connector to Impala.

SAS/ACCESS Interface to Impala supports Impala Server 3.2 or a later version. It also supports the ODBC Driver for Impala, version 2.6.9 or a later version.

---

## Requirements for SAS In-Database Technologies

SAS In-Database Technologies is a technology bundle that is included with multiple SAS Viya product offerings. It includes products that support distributed data sources and have distinct system requirements. To use SAS In-Database Technologies, SAS Embedded Process must be installed.

## Requirements for SAS In-Database Technologies for Hadoop

SAS In-Database Technologies for Hadoop supports the following distributions of Hadoop:

- Amazon Web Services EMR 5.30
- Cloudera CDH 6.2
- Cloudera CDP 7.1

---

**Note:** See the note below for information about Spark support.

---

- Hortonworks HDP 3.1

Execution of SAS Embedded Process for Hadoop on the Spark platform requires Spark 2.

---

**Note:** SAS Embedded Process does not support Hadoop MapReduce when running on Cloudera CDP 7.1. As a result, Spark must be installed and configured on Cloudera CDP 7.1.

---

## Requirements for SAS In-Database Technologies for Teradata

SAS In-Database Technologies for Teradata supports the following products:

- Teradata Vantage Advanced SQL Engine version 16.20 or later
- Teradata CLIV2 client libraries, TTU 16.20 or later for Linux (64-bit libraries)

SAS In-Database Technologies for Teradata also supports Teradata Vantage on the following cloud platforms:

- Amazon Web Services

- Google Cloud Platform
- Microsoft Azure
- Teradata Cloud
- VMware

---

## Requirements for SAS/ACCESS Interface to JDBC

SAS/ACCESS Interface to JDBC includes SAS Data Connector to JDBC. SAS/ACCESS Interface to JDBC enables access to relational databases by means of SQL and the Java Database Connectivity (JDBC) API.

A JDBC driver is required for the data source from which you want to access data. JDBC drivers are available from DBMS vendors and other third-party JDBC driver developers.

---

## Requirements for SAS/ACCESS Interface to Microsoft SQL Server

SAS/ACCESS Interface to Microsoft SQL Server includes SAS Data Connector to Microsoft SQL Server.

SAS/ACCESS Interface to Microsoft SQL Server supports the following products:

- Amazon RDS Microsoft SQL Server (Microsoft SQL Server 2012 or later)
- Google Cloud Platform Cloud SQL for SQL Server version 2017 or later
- Microsoft SQL Server 2012 or later
- Microsoft Azure SQL Database
- Microsoft Azure SQL Database Managed Instance
- Microsoft Azure Synapse

Required client software is included with your SAS Viya installation. Using a Data Source Name (DSN) to connect requires some post-installation configuration of your Kubernetes deployment.

---

## Requirements for SAS/ACCESS Interface to MongoDB

SAS/ACCESS Interface to MongoDB includes SAS Data Connector to MongoDB.

SAS/ACCESS Interface to MongoDB requires the MongoDB C Driver (“libmongoc,” the official client library for C applications). You can obtain the latest MongoDB C driver from the following website: <http://mongoc.org/>. SAS/ACCESS Interface to MongoDB requires the following DBMS products:

- MongoDB version 4.0 or later
- MongoDB C Driver version 1.17

---

## Requirements for SAS/ACCESS Interface to MySQL

SAS/ACCESS Interface to MySQL includes SAS Data Connector to MySQL.

SAS/ACCESS Interface to MySQL requires MySQL Client version 5.6 or later.

The following DBMS products are supported:

- Amazon Aurora (MySQL engine version 5.6 or later)
- Amazon RDS MariaDB (engine version 10.1 or later)
- Amazon RDS MySQL (engine version 5.6 or later)
- Azure Database for MySQL (engine version 5.6 or later)
- Google Cloud SQL for MySQL (engine version 5.6 or later)
- MySQL Server version 5.6 or later
- MariaDB 10.1 or later
- SingleStore (MemSQL) 6.0 or later
- Oracle MySQL Database

---

## Requirements for SAS/ACCESS Interface to Netezza

SAS/ACCESS Interface to Netezza includes SAS Data Connector to Netezza.

SAS/ACCESS Interface to Netezza requires an ODBC driver from IBM. To obtain the appropriate IBM Netezza ODBC driver, contact IBM Technical Support at (877) 426-6006 or visit the IBM Fix Central website: <http://www.ibm.com/support/fixcentral>.

SAS Viya supports the following DBMS products:

- IBM Netezza version 7.2.1 or later
- Client utilities for IBM Netezza version 7.2.1.10 or later

For best results, match the Netezza ODBC client version with the version of the Netezza server where it will be connected. For example, if you have IBM Netezza Interface server release 7.2.1, you should use the ODBC client driver release 7.2.1 with SAS/ACCESS Interface to Netezza.

---

## Requirements for SAS/ACCESS Interface to ODBC

SAS/ACCESS Interface to ODBC includes SAS Data Connector to ODBC. SAS/ACCESS Interface to ODBC enables access to multiple data source types by means of a generic ODBC driver.

Before you can use SAS Viya with ODBC, an ODBC driver is required for the data source from which you want to access data. ODBC drivers are often available from DBMS vendors and other third-party ODBC driver developers. Your ODBC driver must comply with the ODBC 3.5 (or later) specification.

---

**Note:** The ODBC driver that you select might require additional DBMS software in order to enable network access.

---

---

## Requirements for SAS/ACCESS Interface to Oracle

SAS/ACCESS Interface to Oracle (on SAS Viya) includes SAS Data Connector to Oracle.

SAS/ACCESS Interface to Oracle requires the Oracle client 12c or later (64-bit libraries).

SAS Viya requires the following Oracle instances:

- Amazon RDS Oracle 12c or later
- Oracle Cloud Platform 12c or later
- Oracle Database 12c or later

Obtain the path to the volume on the Oracle server to which you want to point your SAS Viya deployment. This information is used later in the deployment.

---

## Requirements for SAS/ACCESS Interface to PC Files

SAS/ACCESS Interface to PC Files includes SAS Data Connector to PC Files.

SAS/ACCESS Interface to PC Files enables access to the following file formats:

- .jmp
- .spss
- .stata
- .xlsx or .xls

Additional Microsoft file formats can be accessed via the PCFILES LIBNAME engine, which is included with SAS/ACCESS Interface to PC Files.

No additional software is required.

---

## Requirements for SAS/ACCESS Interface to the PI System

SAS/ACCESS Interface to the PI System uses the PI System Web API, which is HTTPS-based and RESTful. No PI System client software is required to be installed on the nodes where SAS is running.

However, the PI System Web API (PI Web API 2019 or later) must be installed and activated on the host machine from which the user connects.

---

## Requirements for SAS/ACCESS Interface to PostgreSQL

SAS/ACCESS Interface to PostgreSQL (on SAS Viya) includes SAS Data Connector to PostgreSQL.

SAS/ACCESS Interface to PostgreSQL supports:

- Amazon Aurora (PostgreSQL engine version 9.6 or later)
- Amazon RDS PostgreSQL (engine version 9.6 or later)
- Azure Database for PostgreSQL (engine version 9.6 or later)
- EnterpriseDB PostgreSQL version 9.6 or later
- Google Cloud Platform Cloud SQL for PostgreSQL (engine version 9.6 or later)
- PostgreSQL Database version 9.6 or later

Required client software is included with your SAS Viya installation. Using a Data Source Name (DSN) to connect requires some post-installation configuration of your Kubernetes deployment.

---

## Requirements for SAS/ACCESS Interface to Salesforce

SAS/ACCESS Interface to Salesforce includes SAS Data Connector to Salesforce.

SAS/ACCESS Interface to Salesforce requires a Salesforce user account that has API access enabled. SAS/ACCESS Interface to Salesforce supports Salesforce API access, version 46.0 or later.

---

## Requirements for SAS/ACCESS Interface to SAP ASE

SAS/ACCESS Interface to SAP ASE requires SAP ASE (formerly Sybase) Open Client SDK, Release 15.7 or later (64-bit libraries).

The database administrator must install two SAP ASE (Sybase) stored procedures on the target SAP server. These files are available in a compressed TGZ archive for download from the SAS Support site at <https://support.sas.com/downloads/package.htm?pid=2458>.

---

## Requirements for SAS/ACCESS Interface to SAP HANA

SAS/ACCESS Interface to SAP HANA includes SAS Data Connector to SAP HANA.

SAS/ACCESS Interface to SAP HANA supports SAP HANA SPS 11 Server or a later version and requires SAP HANA ODBC Client (64-bit) for SPS 11 or later.

In addition, the SAP HANA ODBC Client requires a compatible ODBC Driver Manager, such as the Unix ODBC Driver Manager.

---

## Requirements for SAS/ACCESS Interface to R/3

SAS/ACCESS Interface to SAP R/3 supports SAP NetWeaver 7.0 (Application Server ABAP) or later and requires the 64-bit SAP NetWeaver RFC Library, Release 7.20 or later, which is provided by SAP AG.

---

## Requirements for SAS/ACCESS Interface to Snowflake

SAS/ACCESS Interface to Snowflake includes SAS Data Connector to Snowflake.

SAS/ACCESS Interface to Snowflake requires the Snowflake ODBC driver version 2.21.5 or later.

In addition, the Snowflake ODBC Client requires a compatible ODBC Driver Manager, such as the Unix ODBC Driver Manager.

---

## Requirements for SAS/ACCESS Interface to Spark

SAS/ACCESS Interface to Spark includes SAS Data Connector to Spark.

SAS/ACCESS Interface to Spark supports Spark server 2.4 and requires Spark client JAR files from the Hadoop cluster where the Spark server is running.

SAS/ACCESS Interface to Spark supports one Hadoop distribution: Hortonworks HDP 3.1.

---

## Requirements for SAS/ACCESS Interface to Teradata

SAS/ACCESS Interface to Teradata includes SAS Data Connector to Teradata.

SAS/ACCESS Interface to Teradata requires Teradata CLIV2 client libraries, TTU 16.20 or later, and supports Teradata Vantage SQL Engine version 16.20 in addition to the following platforms:

- Teradata IntelliCloud
- Teradata Vantage on Amazon Web Services
- Teradata Vantage on Google Cloud Platform
- Teradata Vantage on Microsoft Azure

---

## Requirements for SAS/ACCESS Interface to Vertica

SAS/ACCESS Interface to Vertica includes SAS Data Connector to Vertica.

SAS/ACCESS Interface to Vertica requires Vertica ODBC Client version 9.1 or later and supports Vertica Analytic Database version 9.1 or later.

To obtain the Vertica Client ODBC driver, contact your database administrator or visit the Vertica website: <https://my.vertica.com/download/vertica/client-drivers>.

---

## Requirements for SAS/ACCESS Interface to Yellowbrick

SAS/ACCESS Interface to Yellowbrick includes the required PostgreSQL ODBC driver.

Yellowbrick Database version 4.0.0-23452 or later is supported.

---

# Security Requirements

---

## About SAS Viya Security Features

SAS Viya provides Transport Layer Security (TLS) for encryption of data in motion and supports the Advanced Encryption Standard (AES) for encryption of data at rest. With minimal configuration, you can deploy and use the SAS Security Certificate Framework, which contains tools that leverage Kubernetes features to provide encryption for SAS applications.

Multiple methods are supported for user authorization and authentication. Most user account configuration occurs after the deployment has completed. Single sign-on and multi-factor authentication are also supported.

## Encryption Overview

The deployment enables TLS 1.2 for connections to the cluster, from SAS Viya components to your IT infrastructure, and among the pods. The SAS Security Certificate Framework uses cert-manager, which is an open-source Kubernetes tool, to manage CA certificates. By default, cert-manager generates the CA-signed certificates that it then distributes; however, you can provide your own certificates if desired. Custom certificates must be in PEM format.

The CAS server supports encryption for tables in caslibs. SAS Viya uses AES with 256-bit keys to encrypt stored data. The encryption applies to source tables, not to tables that are resident in memory. Encryption can be applied to individual tables or to all tables in a library. Each table can have a unique encryption key, or a single key can be set at the library level in order to have a shared key for all tables in the library.

The default deployment enables TLS for an NGINX Ingress (“full TLS”). SAS provides customizations and instructions that configure SAS services to define Ingress objects that connect into an ingress-nginx controller.

For more information about TLS, see “[Encryption Overview](#)” in *Encryption in SAS Viya: Data in Motion* for relevant conceptual information and procedures.

## SAS Security Certificate Framework Requirements

To apply TLS encryption to your SAS Viya deployment, an add-on is required in order to enable Kubernetes to manage and distribute certificates. Install cert-manager in the cluster before you start the deployment. See “[Kubernetes Cluster Requirements](#)” on [page 3](#) for details.

In addition, cert-manager requires an issuer named sas-viya-issuer. The command to create this issuer is included by default in the kustomization.yaml file.

Configure cert-manager to automatically delete secrets when they are no longer being used. For instructions, see “[Delete Secrets That Are No Longer Used](#)” in *Encryption in SAS Viya: Data in Motion*.

Instructions for customizing TLS for your cluster are provided in the README file at `$deploy/sas-bases/examples/security/README.md` (for Markdown format) or at `$deploy/sas-bases/docs/configure_network_security_and_encryption_using_sas_security_certificate_framework.htm` (for HTML format).

## Pod Security Policies

If your Kubernetes cluster has pod security policies enabled, a conflict with a SAS Viya seccomp annotation can cause the deployment to fail. SAS Viya requires `runtime/default` to be allowed in the environment. You can check pod security policies by running the following command on your cluster:

```
kubect1 get psp restricted -o custom-columns=NAME:.metadata.name,
"SECCOMP":".metadata.annotations.seccomp\.security\.alpha\.kubernetes\.io/
allowedProfileNames"
```

You can resolve this conflict by updating a variable in the pod security policy. For more information, see the following SAS Note: <https://support.sas.com/kb/67/349.html>.

---

## Identity Provider and Authentication Requirements

SAS Viya supports LDAP and the System for Cross-domain Identity Management (SCIM) for user and group identities. Multiple identity providers are supported, including Microsoft Azure Active Directory.

## Supported Authentication Methods

SAS Viya supports the following authentication methods:

- LDAP, Kerberos, and single sign-on are supported for users who are signing in to the environment.
- Single sign-on using Security Assertion Markup Language (SAML) or OpenID Connect. Single sign-on support is limited to web log-ins.

## LDAP Requirements

When the deployment completes, SAS Viya is configured by default to use LDAP. The following requirements apply to your LDAP server:

- SAS Viya must have Read access to your LDAP server.
- In order to bind to the LDAP server, SAS Viya requires either a system account (with a userDN and password) or anonymous binding.
- If the mail attribute is specified for LDAP accounts, it must have a non-null value that is unique for each user.
- LDAPS is supported, but the required certificates are not configured automatically by the deployment process.

You will configure SAS Viya identities as a post-deployment step. Instructions for setting up identities with LDAP are provided in [Identity Management](#).

## SCIM Requirements

SAS Viya supports SCIM 2.0.

When the deployment completes, SAS Viya is configured by default to use LDAP. If you intend to use SCIM, you must disable LDAP by logging in to SAS Environment Manager after the deployment process has completed.

Single sign-on using SAML or OpenID Connect (OIDC) is required if you configure a SCIM identity provider.

With SCIM, you can use Okta or Microsoft Azure Active Directory to populate user and group identities in SAS Viya. In a GCP environment, Google Identity is not yet supported for user provisioning; use Okta or Azure Active Directory instead. SCIM requires network access to the SAS Viya deployment using HTTPS, and it requires a certificate signed by a public certificate authority. The SCIM IdP also requires a long-lived token to access the SAS Viya APIs.

Check your firewall settings so that IP addresses used by the SCIM IdP are allowed to reach the SAS Viya Ingress. In Microsoft Azure, you can configure your network security group to allow Azure Active Directory to communicate with resources in your virtual network by enabling inbound access for the AzureActiveDirectory service tag.

You configure SAS Viya identities as a post-deployment step. Instructions for setting up identities with SCIM are provided in [Identity Management](#).

## About Roles and Permissions

The Kubernetes specification accepts system-delineated roles and user-facing roles. SAS Viya applies a *least-privileges* model that grants to each role the minimum access that is required to deploy and run the application. SAS Viya resources are controlled by labels that describe installation requirements and that determine the privileges for enabling access to these resources.

To deploy SAS Viya, an administrator with cluster permissions must run the commands to apply the application resources to the cluster. The requirement to use label selectors applies to `kubectl` commands that require the highest level of access: `kubectl apply` and `kubectl delete`. For `kubectl apply`, the label selector explicitly avoids race conditions. Typical race conditions occur when the Kubernetes API processes objects in parallel that have dependencies on each other.

Here is an example of a `kubectl` command that uses a label selector:

```
kustomize build | kubectl apply --kubeconfig=cluster-admin.conf --selector="sas.com/admin=cluster-wide" -f -
```

The selector `sas.com/admin=cluster-wide` indicates that the command can be executed only by a user with the default cluster-admin ClusterRole. This role grants the equivalent of super-user privileges.

To provide granular privileges and to ensure that the least-privileges model is enforced, SAS applies a custom resource label to one category of SAS Viya resources. The following table summarizes SAS Viya resources and their corresponding labels:

**Table 10** SAS Viya Label Selectors

Label	Description	SAS Viya Resources
cluster-wide	This label applies to resource modifications that affect the entire cluster. This category includes global objects, such as CRDs or ClusterRoles, that could affect deployments in other namespaces. Interacting with resources in this category also requires cluster-admin permissions. These resources are defined under <a href="#">Cluster APIs</a> in the Kubernetes API Reference documentation.	CustomResourceDefinition; ServiceAccount; ClusterRole; Role; PersistentVolume
cluster-local	This is a SAS custom label for resources that require cluster-admin permissions, but it limits the impact of changes to the namespace, with a reduced impact on the cluster. These resources are defined under <a href="#">Config and Storage APIs</a> and	ResourceQuota; LimitRange; Secret; ConfigMap; RoleBinding; ClusterRoleBinding; PersistentVolumeClaim; PodTemplate

Label	Description	SAS Viya Resources
	<a href="#">Metadata APIs</a> in the Kubernetes API Reference.	
namespace	<p>This label applies to the following resource types:</p> <ul style="list-style-type: none"> <li>resources for managing and running containers</li> <li>resources that provide external access to workloads by means of a LoadBalancer or NodePort service</li> </ul> <p>These resources are defined under <a href="#">Workloads APIs</a> and <a href="#">Service APIs</a> in the Kubernetes API Reference.</p>	Service; Deployment; ReplicaSet; ReplicationController; Job; CronJob; StatefulSet; DaemonSet

Some Kubernetes Role-Based Access Control (RBAC) policies are enabled by default. Because objects of each kind are deployed, permissions for each resource should be granted to the appropriate role. During RBAC planning for your SAS Viya deployment, consider that some user-facing ClusterRoles need to allow admin users to include rules for custom resources.

## User Accounts

You can set up SAS Viya users and groups with SAS Environment Manager after your software has been deployed.

Container images are delivered with a critical process that is owned by the UID 1001. This UID cannot be changed. Verify that no user accounts in your LDAP directory are using this UID.

An account that is named `sasboot` is also created during the deployment and has an expired password. It is an administrator account that is used for preliminary logon to SAS Environment Manager. It is internal only to SAS and does not exist in your identity provider. For more information, see [Sign In as the sasboot User](#).

## Service Accounts

During the deployment process, required infrastructure services are automatically started in separate containers within the SAS Viya namespace. Some of these services are owned by a Kubernetes service account, which is created by the deployment process. Each service account has a very limited role. The use of multiple service accounts to run services makes it easier to secure your SAS Viya environment because each service account has only the specific privileges that it requires to run one service and has no additional privileges.

The following Kubernetes service accounts are created by the deployment process:

Table 11 Required Kubernetes Service Accounts

Service Account	Description
postgres-operator	Owner of the PostgreSQL operator. Supports either an internal or an external instance of PostgreSQL for SAS Infrastructure Data Server.
pgo-backrest	Owner of the Crunchy Data service to support an internal PostgreSQL instance for SAS Infrastructure Data Server. All the containers that are launched by the PostgreSQL operator as part of the pgbbackup service use this service account.  Crunchy Data service role bindings are part of the custom resource definition for the PostgreSQL operator. This account and container are not present if you selected to deploy with an external instance of PostgreSQL.
pgo-default	Used by the sas-crunchy-data-postgres-backrest-shared-repo service.
pgo-pg	Used by the Crunchy Data service, sas-crunchy-data-postgres.
pgo-target	Owner of the Crunchy Data service account to support an internal PostgreSQL instance for SAS Infrastructure Data Server. Crunchy Data service role bindings are part of the custom resource definition for the PostgreSQL operator. This account and container are not present if you chose to deploy with an external instance of PostgreSQL.
sas-rabbitmq-server	Owner of the SAS Message Broker service.
sas-cas-operator	Owner of the CAS server operator.
sas-cas-server	Owner of the CAS server pods. It enables CAS server pod security policies.  This account has both a role and a role binding of sas-cas-server.
sas-config-reconciler	Owner of the service that keeps pods synchronized with configuration data, relaunching them when their configuration changes.
sas-consul-server	Owner of the Consul key-value store.  SAS Configuration Server is based on HashiCorp Consul.
sas-data-server-utility	Owner of a required account to support SAS Infrastructure Data Server. Manages database creation, Consul registration, and PostgreSQL validation. It runs in a container that is separate from the PostgreSQL Operator. This container is always present, regardless of whether you deployed with an internal or external PostgreSQL instance.

Service Account	Description
sas-launcher	Owner of a service that launches Kubernetes jobs and checks the status of those jobs.
sas-prepull	A service account that improves start-up times by pre-staging the sas-programming image to nodes that require it.
sas-viya-backuprunner	A service account that is used for backup operations.

Additional service accounts might be present if you deployed SAS Viya products that require them.

## SAS Infrastructure Data Server Requirements

### PostgreSQL Deployment Options

SAS Infrastructure Data Server stores SAS Viya user content, such as reports, authorization rules, selected source definitions, attachments, audit records, and user preferences. SAS Viya uses High Availability (HA) PostgreSQL for SAS Infrastructure Data Server. Take some time to consider your options for deploying a PostgreSQL instance to support the data server.

### Internal versus External PostgreSQL Instances

An instance of PostgreSQL is required for the SAS Data Infrastructure Server. You can allow SAS to automatically deploy an instance of PostgreSQL, which is described as an *internal* instance, or you can provide your own PostgreSQL, which is an *external* instance. Before deploying your software, you must decide which type of instance you want.

**IMPORTANT** After your SAS Viya software has been deployed, you cannot change the PostgreSQL deployment type without redeploying the entire software package.

In addition to the following requirements, both types of PostgreSQL instance require persistent storage and PVCs. For more information, see [“Persistent Storage Volumes, PersistentVolumeClaims, and Storage Classes” on page 8](#).

### Internal PostgreSQL

If you decide to deploy an internal instance of PostgreSQL for SAS Infrastructure Data Server, SAS configures and maintains the deployment for you.

SAS deploys a comprehensive PostgreSQL container operator suite provided by [Crunchy Data](#) for this purpose.

## Requirements for External PostgreSQL

If you decide to deploy an external instance of PostgreSQL, you are responsible for configuring and maintaining the PostgreSQL deployment. Providing your own instance enables you to restrict the ability of SAS services to access the database when a single PostgreSQL instance houses multiple databases. The following requirements apply:

- PostgreSQL versions 11 and 12 are the only supported versions of PostgreSQL for SAS Viya.
- Managed PostgreSQL instances that are hosted by your cloud provider, such as Microsoft Azure Database for PostgreSQL, Amazon RDS for PostgreSQL, or Cloud SQL for PostgreSQL, are supported.

---

**Note:** SAS Viya is supported on AWS and on GCP only for the Early Adopter program at this time.

---

- All PostgreSQL extensions must be installed.
- The database user must have standard permissions to Create, Read, Update, and Delete (CRUD) data.
- The external PostgreSQL instance must already be running and the database user must be created before you start the SAS Viya deployment.
- A low-latency, high-bandwidth environment is required.

Placing the external PostgreSQL database in a separate data center, region, or availability zone from the rest of your SAS Viya components might lead to increased latency and reduced bandwidth. Such conditions are likely to cause a degraded overall performance of the environment. Before attempting to deploy SAS Viya with an external PostgreSQL database, it is important to test with it in order to balance cost and performance considerations and to confirm that performance is not adversely affected.

- (GCP only) You must use Google's Cloud SQL Proxy to access your Cloud SQL for PostgreSQL server.

Additional steps are required in order to support full TLS. For more information, see the README file titled "Configure an External PostgreSQL Instance for SAS Viya" at `$deploy/sas-bases/overlays/configure-postgres/README.md` (for Markdown format) or at `$deploy/sas-bases/docs/configure_an_external_postgresql_instance_for_sas_viya.htm` (for HTML format).

# Client Requirements

## Web Browsers

End users can access the product user interfaces for SAS Viya applications from a desktop computer, using a supported web browser. Because SAS software is not installed on this machine, the requirements are minimal. UNIX and 64-bit Windows operating systems are supported.

Your browser must be enabled for JavaScript.

Some SAS Viya user interfaces include some advanced features that require recent versions of popular web browsers. Client computers that access these interfaces should adhere to the following guidelines:

*Table 12 Supported client operating systems and browsers*

Operating System	Web Browser (64-bit)	Web Browser (32-bit)
Apple OS X 10.13 and later	Apple Safari 13.0	
Linux 64-bit (x86-64)	Google Chrome 77.0 and later Mozilla Firefox 68.0 and later	Google Chrome 77.0 and later
Microsoft Windows 10 (64-bit)	Google Chrome 77.0 and later Mozilla Firefox 68.0 and later Microsoft Edge on Chromium 80.0 and later	Google Chrome 77.0 and later Mozilla Firefox 68.0 and later Microsoft Edge on Chromium 80.0 and later
Microsoft Windows 10 (32-bit)		Google Chrome 77.0 and later Mozilla Firefox 68.0 and later Microsoft Edge on Chromium 80.0 and later
Microsoft Windows 8.1 (64-bit)	Google Chrome 77.0 and later Mozilla Firefox 68.0 and later	Google Chrome 77.0 and later Mozilla Firefox 68.0 and later
Microsoft Windows 8.1 (32-bit)		Google Chrome 77.0 and later Mozilla Firefox 68.0 and later
Microsoft Windows 8 (64-bit)	Google Chrome 77.0 and later	Google Chrome 77.0 and later Mozilla Firefox 68.0 and later
Microsoft Windows 8 (32-bit)		Google Chrome 77.0 and later

Operating System	Web Browser (64-bit)	Web Browser (32-bit)
		Mozilla Firefox 68.0 and later
Microsoft Windows 7 (64-bit)	Google Chrome 77.0 and later	Google Chrome 77.0 and later Mozilla Firefox 68.0 and later
Microsoft Windows 7 (32-bit)		Google Chrome 77.0 and later Mozilla Firefox 68.0 and later
Microsoft Windows Server 2012 R2		Google Chrome 77.0 and later Mozilla Firefox 68.0 and later
Microsoft Windows Server 2012		Google Chrome 77.0 and later Mozilla Firefox 68.0 and later
Microsoft Windows Server 2008 R2		Google Chrome 77.0 and later Mozilla Firefox 68.0 and later

## Mobile Platform Support

Some SAS Viya user interfaces are not currently supported on mobile devices.

When SAS Visual Analytics Apps for iOS, Android, and Windows 10 are installed on a mobile device, users of the following interfaces can view and explore reports on the mobile device:

- SAS Visual Analytics
- SAS Visual Statistics
- SAS Visual Machine Learning

SAS Visual Analytics Apps are supported on iOS and Android smartphones and tablets, and on computers and tablets running Microsoft Windows 10. Each app is written specifically for its operating system and provides native support for these devices so that you can view and explore reports on your mobile device using a touchscreen. You can download the apps for free from the Apple App Store, Google Play, or the Microsoft Windows Store.

## Requirements for SAS® for Microsoft® 365 Clients

SAS for Microsoft 365 enables SAS analytics to access reports directly from Microsoft Excel 365 and provides integrated features. Some requirements for SAS for Microsoft 365 differ from those of other SAS Viya offerings.

## Supported Browsers for the Web Application

After you deploy SAS Viya and configure SAS for Microsoft 365, the SAS for Microsoft 365 web application is available from the Office 365 Excel web client in a supported web browser. When it runs in a browser, the SAS for Microsoft 365 web application supports only the following web browsers:

- Google Chrome 77.0 and later
- Microsoft Edge on Chromium 80.0 and later
- Mozilla Firefox 68.0 and later
- Apple Safari 13.0 and later on macOS

Browsers running on mobile or touchscreen devices are not supported at this time.

## Operating Environments for the Desktop Application

SAS for Microsoft 365 can also run as an installed add-in to a desktop version of Microsoft Excel. In order to deploy the SAS for Microsoft 365 installed add-in for Microsoft Excel, use a machine that meets the following requirements:

- Microsoft Windows 10 or macOS
- Microsoft Excel: Microsoft 365 version 16.0.13530.20424 or later
- Microsoft Excel requires Microsoft Edge on Chromium with WebView2

For more information, see [this Microsoft article](#).

## Security Requirements

Microsoft requires add-ins, such as SAS for Microsoft 365, to run in an iframe in the Office 365 web application. To ensure that SAS for Microsoft 365 works properly, your SAS administrator must update these properties in SAS Environment Manager after deployment. After changing properties in SAS Environment Manager, you must restart the SAS Viya pods for the changes to take effect.

For more information and step-by-step instructions, see “[Steps for SAS Administrator](#)” in *SAS for Microsoft 365: User's Guide*.

---

## Screen Resolution

The minimum screen resolution for each client machine that will access the SAS Viya user interfaces is 1280 x 1024.

---

# Verify the Environment

---

## Run a Pre-installation Check

SAS recommends that you run the SAS Pre-Install Check Utility before deploying the software.

SAS Pre-Install Check retrieves information about your Kubernetes cluster and verifies that the cluster meets the documented requirements to deploy SAS Viya. The utility generates a report.html that displays the results of the checks that it performs.

The files that are required to create the Pre-Install Check Report and the instructions for running the utility are available at the [SAS GitHub site](#). The README file contains links to “Pre-installation of SAS Viya System Requirements” and other tools that you can use once the deployment has completed.