



# SAS<sup>®</sup> Viya<sup>®</sup> 4 Architecture


Prepared for SUGA

June 2021




# Why SAS Viya?

- Deep Learning
- Scalable Speed
- Cloud
- High Availability
- Support Open Source
- Democratization of Analytics



# SAS Viya... Becoming “Cloud Native”

- Cloud Portability
- Ease of Management
- Simplified Licensing
- Continuous Delivery
- Cloud Service Integration
- Self-Service Portal
- SaaS



# Why Kubernetes?

## Benefits to SAS and for its customers

### For SAS:

- Enables “containers as currency”
- Provides cloud portability
- Delivers cloud-native on-demand resource allocation
- Fits into our IT & SAS Cloud strategy

### For customers:

- Simplifies s/w management
- Abstracts away complexity of underlying infrastructure
- Industry standard
- Choice of:
  - Managed service on public cloud
  - Commercial
  - Open source

# SAS Viya deployed in Kubernetes

## Microservices and Web Apps



STATELESS NODE POOL

## SAS Compute Services




COMPUTE NODE POOL

## Commodity Services



STATEFUL NODE POOL

## CAS Cluster




CAS NODE POOL


## Ingress

SYSTEM NODE POOL

## Monitoring + Logs



OPS4VIYA



## Microservices and Web Apps




STATELESS NODE POOL

## SAS Compute Services



COMPUTE NODE POOL

## CAS Cluster




CAS NODE POOL



Azure Kubernetes Service (AKS)

-  Azure AD
-  DNS
-  Load Balancer
-  Monitor
-  Service Bus
-  Cosmos DB
-  DB for PostgreSQL


## Microservices and Web Apps



## SAS Compute Services



## CAS Cluster



Elastic Kubernetes Service (EKS)



AWS IAM



Route 53 DNS



Elastic Load Balancing



CloudWatch



Simple Queue Service




Dynamo DB



AWS RDS for PostgreSQL


## Microservices and Web Apps










## SAS Compute Services



## CAS Cluster



Google Kubernetes Engine (GKE)

-  Cloud IAM
-  Cloud DNS
-  Cloud Load Balancing
-  Monitoring
-  Cloud Pub/Sub
-  Cloud Datastore
-  Cloud SQL for PostgreSQL



On-Premises  
Target - August 2021

Microservices and Web Apps



STATELESS NODE POOL

SAS Compute Services



STATEFUL NODE POOL

CAS Cluster



CAS NODE POOL


Ingress

SYSTEM NODE POOL

Monitoring + Logs



OPS4VIYA



# Continuous Delivery Cadences

**2021.1**

## LONG TERM SUPPORT (LTS)

- Available every 6 months
- Seasoned code (Stable-2)
- Update-in-place

**2021.1.1**

## STABLE

- Monthly feature updates
- Includes bug fixes
- Update-in-place

Support for current plus 3 prior releases

LTS ↔ Stable

# Other 2021 Viya Highlights

New Cloud Environments



AWS, GCP, *OpenShift* on VMWare, ...

Deployment Automation



Quickstart provisioning (IaC); Lifecycle management

Personal CAS Server



Single-user CAS Servers on-demand

Kerberos



End-to-end security

# Workload Management

Bringing SAS Grid to SAS Viya 4

# Workload Management in SAS Viya - Roadmap

## SAS Viya 4 – Now

- High availability
- Scalability
- Batch submission of jobs
  - *SASGSUB* → *sas-viya batch*
- Job flow definition/scheduling



# Workload Management in SAS Viya - Roadmap

SAS Viya 4 - LTS 2021.2

## 2021 R&D Priorities – “Parity” and “Migration”

- Option sets for users/groups/roles
- Resource definitions: Jobs  $\leftrightarrow$  Kubernetes nodes
- Queues
- Prioritization
- Preempt/resume
- SAS program checkpoint/restart
- Job Flow Scheduler  $\approx$  Platform Suite for SAS in 9.4



# Workload Management in SAS Viya - Roadmap

SAS Viya 4 - 2022+

## “Eating our own SAS analytics cooking”

- Optimization of cloud resources
  - Cost thresholds
  - Performance SLAs
- Static profiling
- Historic profiling



Questions?