



Alex Daehnrich, Doug Haigh

SAS Viya Workload Management - The next generation of elasticity!



SAS Viya Workload Management-The next generation of elasticity!

Alex Daehnrich
Senior Manager
SAS Workload Management



Alex Daehnrich is leading multiple engineering teams for the SAS solutions under the Workload Management and Orchestration department. His primary role is to drive software development efforts within SAS R&D. For the last 23+ years and prior to holding his current position, he held multiple leadership positions in Consulting, Pre-Sales, Partner Enablement and Research & Development.

Doug Haigh
Distinguished Software Developer
SAS Workload Management




Doug joined SAS in 2005 and has been the primary designer and developer of SAS Grid Manager, SAS/CONNECT and SAS/SECURE products. Doug is currently the principal architect and developer for SAS Workload Orchestrator for v9 and Viya 4.

AGENDA



Motivation



Capabilities

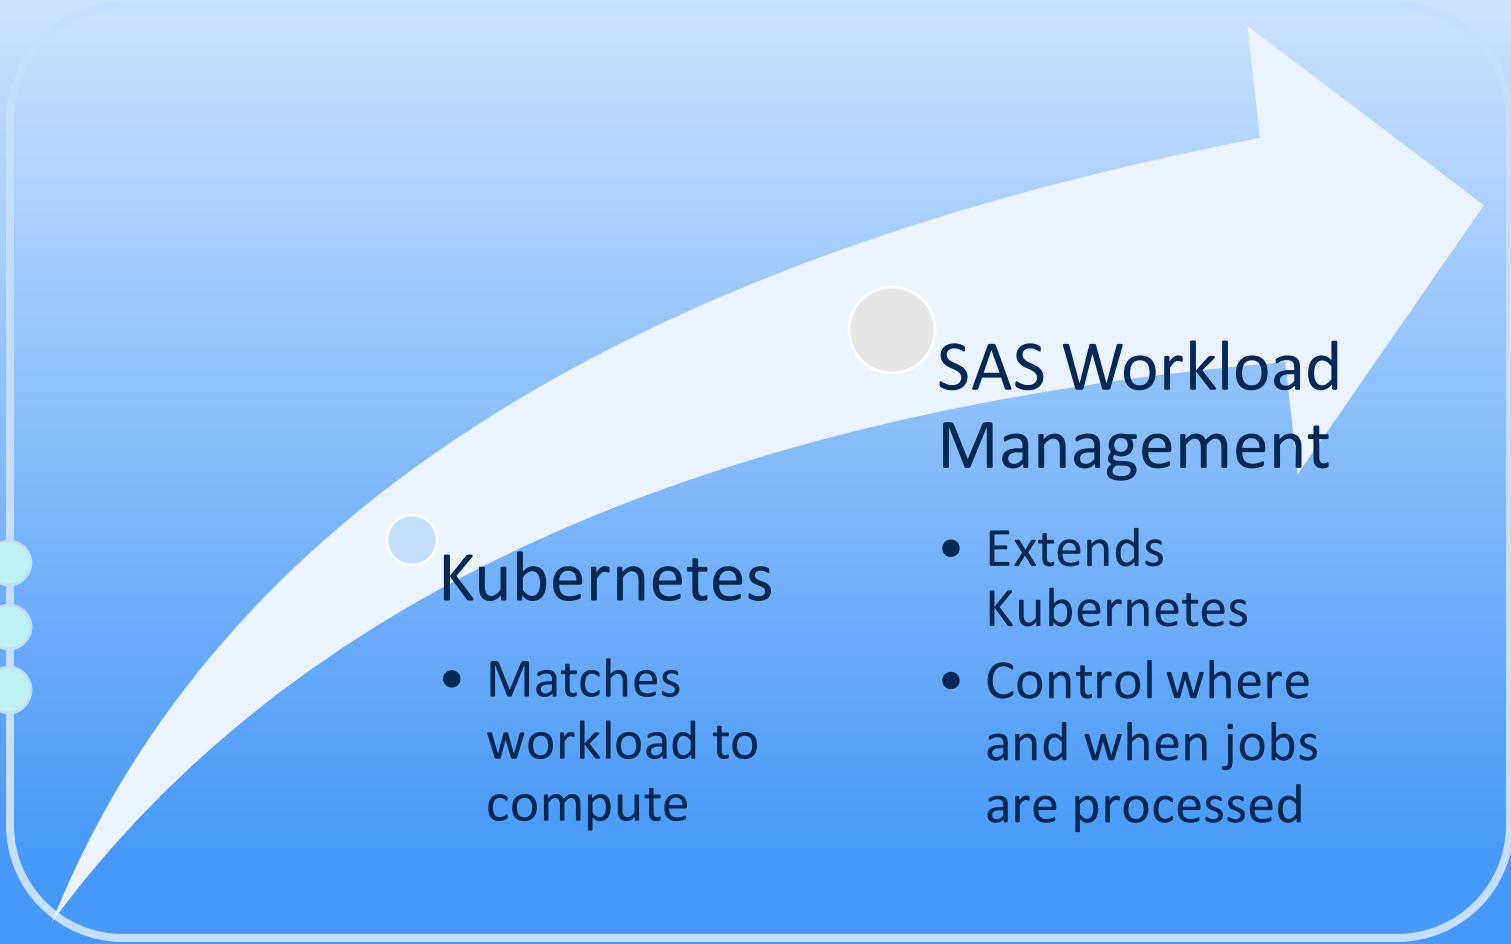



Demo



Future





Kubernetes

- Matches workload to compute

SAS Workload Management

- Extends Kubernetes
- Control where and when jobs are processed

Why Workload Management in Viya?

Is Kubernetes no good?



SAS still spawns servers for individual users



Different groups have different budgets



Different users need different resources



Different amounts of resources may be available at different times



Different users need different limits




Different resources may have different limits

AGENDA



Motivation



Capabilities



Demo



Future



Issues Workload Management Tries to Solve

Allowing SAS Administrator to manage resources as opposed to a Kubernetes administrator



amount of resources



type of resources



Timing – **“when”**



jobs of a certain type



Scheduling resources

Workload Management Benefits

Full Control



USER

- run designated workloads at designated times



JOBS

- limited per user, per host, per queue
- resource limited (run time, memory used, ...)
- prioritized possibly preempting other specific jobs



HOST

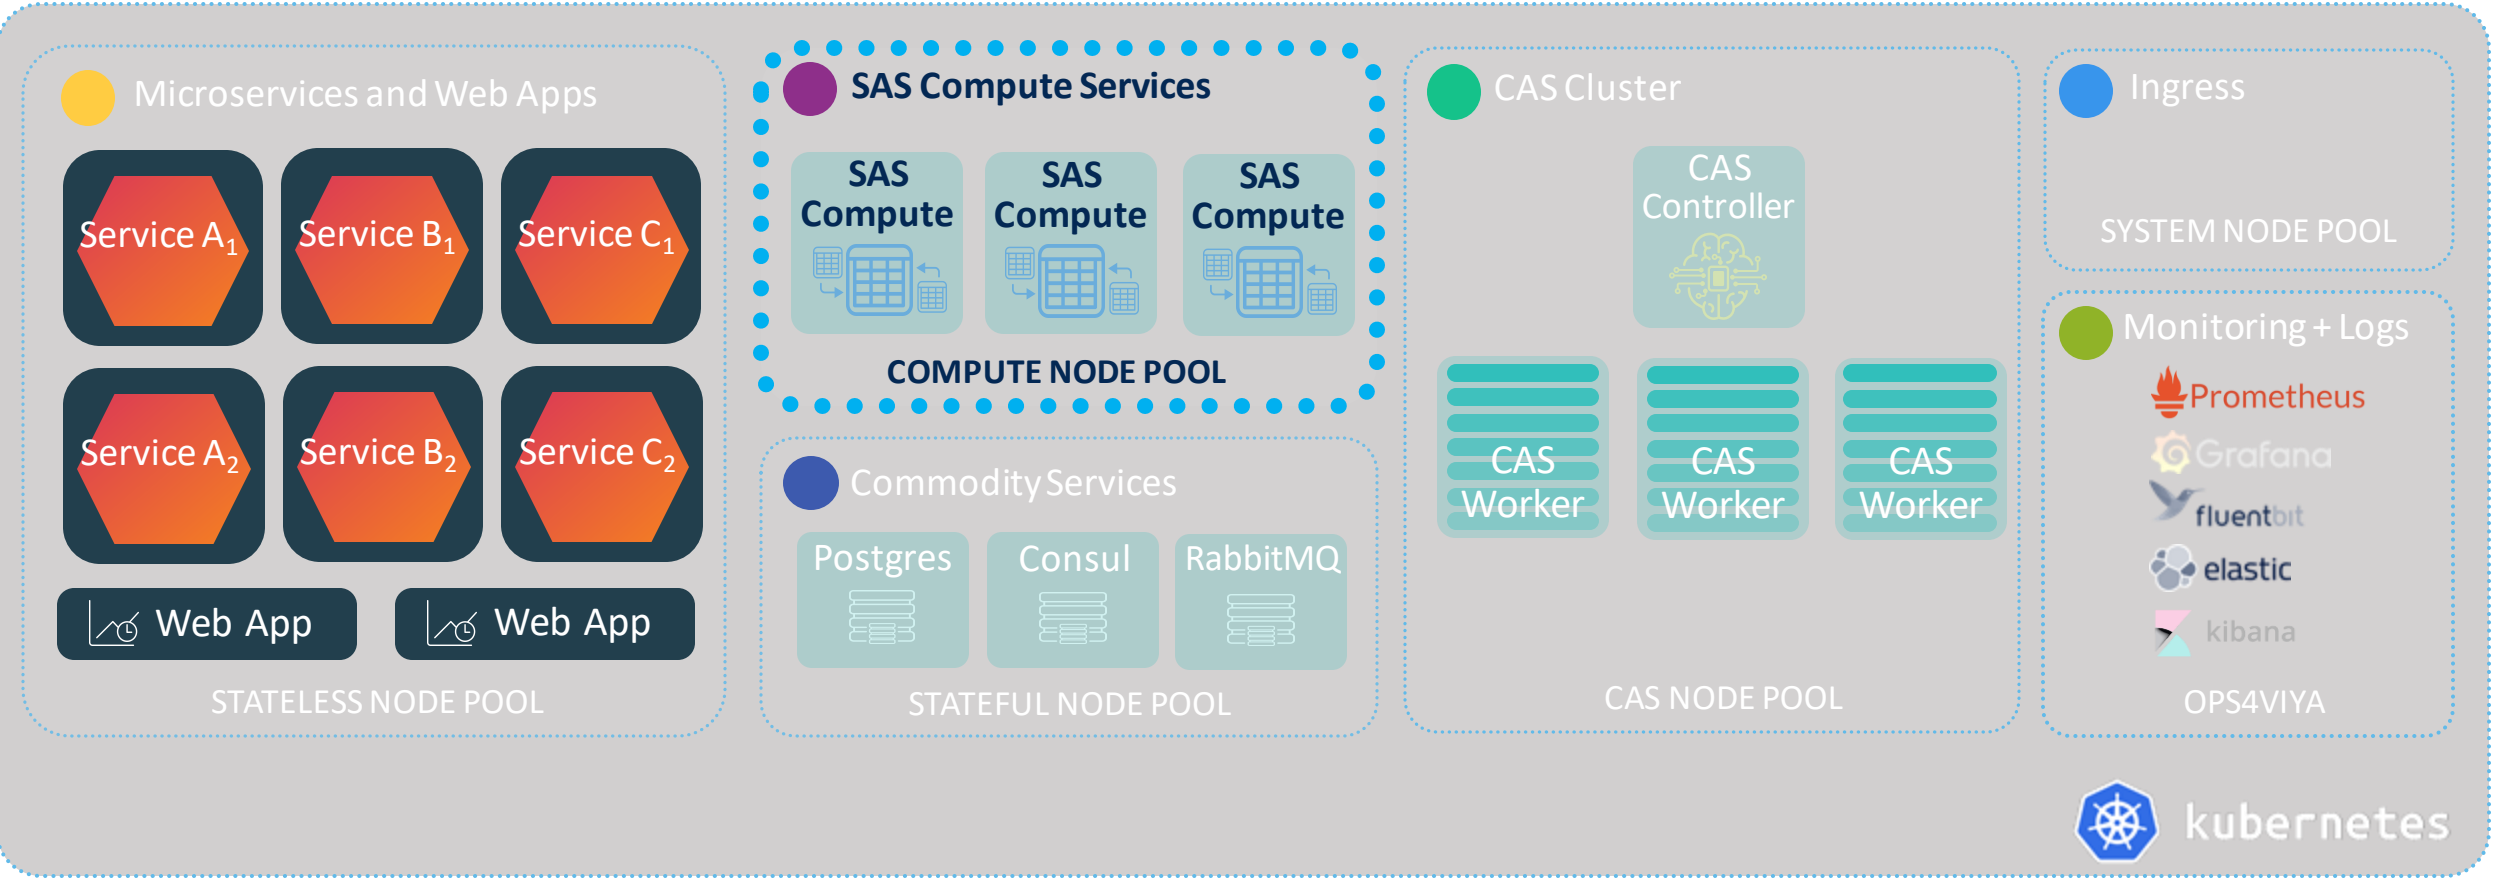
- scheduling and suspension thresholds
- limiting of number of jobs executing at once
- Jobs would go to least utilized host of allowed list of hosts



CONFIGURATION

- Single point of administration

SAS Workload Management is focused on SAS Compute Services



The Solution

How are we solving the problem?

- Implementing existing and **proven queue and policy management** methodologies within SAS Viya with SAS Workload Management
- Full utilization out of the box cloud and Kubernetes capabilities and **extending the functionally** to flexible customize analytical workload execution
- With WLM **adding layer of resilience** to platform
- Enabling **SAS Admins to control** job execution & configuration
- **Providing a new form of elasticity to SAS Customers with WLM Autoscaling**

Balance cost with agility

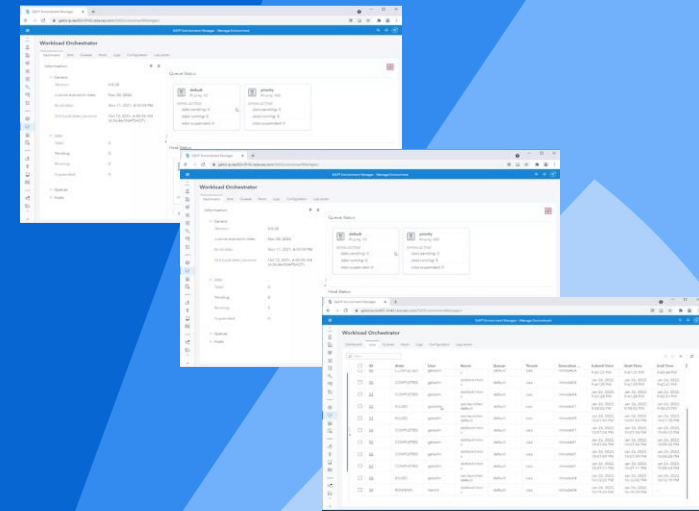
- Queue prioritization
- Multi-tenant support
- Load balancing
- Runs in Kubernetes
- Limit user disruption
- Launcher Context per application
- WLM policy driven autoscaling

Improve throughput, availability and productivity

- Target optimal queue
- Simultaneous job execution
- Workload preemption
- Automatic restart
- Expanded high-availability

Simplify administration


- SAS Workloads
- Configuration
- Monitoring



AGENDA



Motivation



Capabilities



Demo



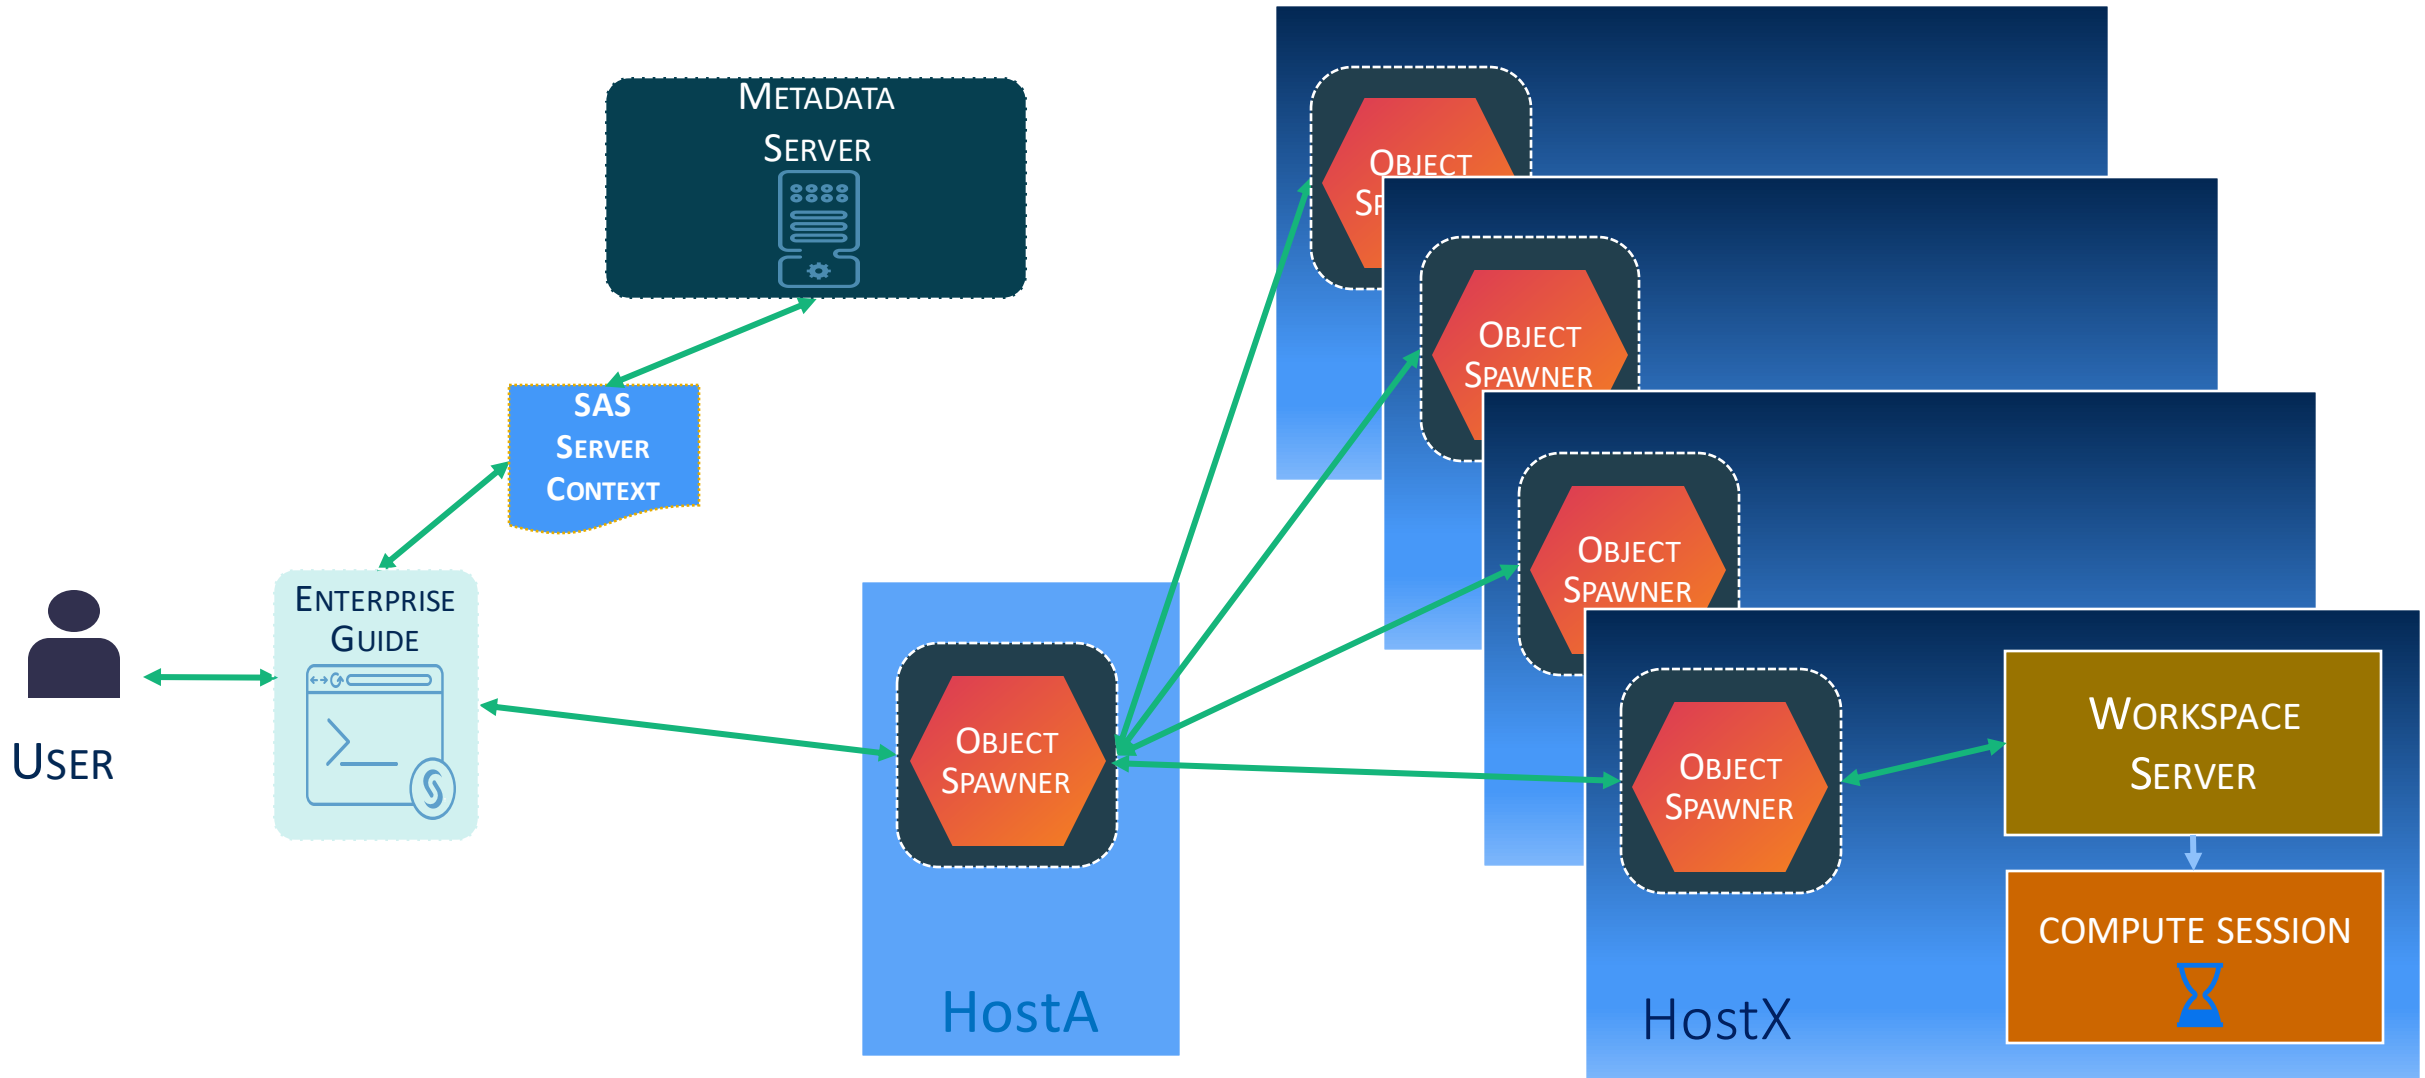
Future



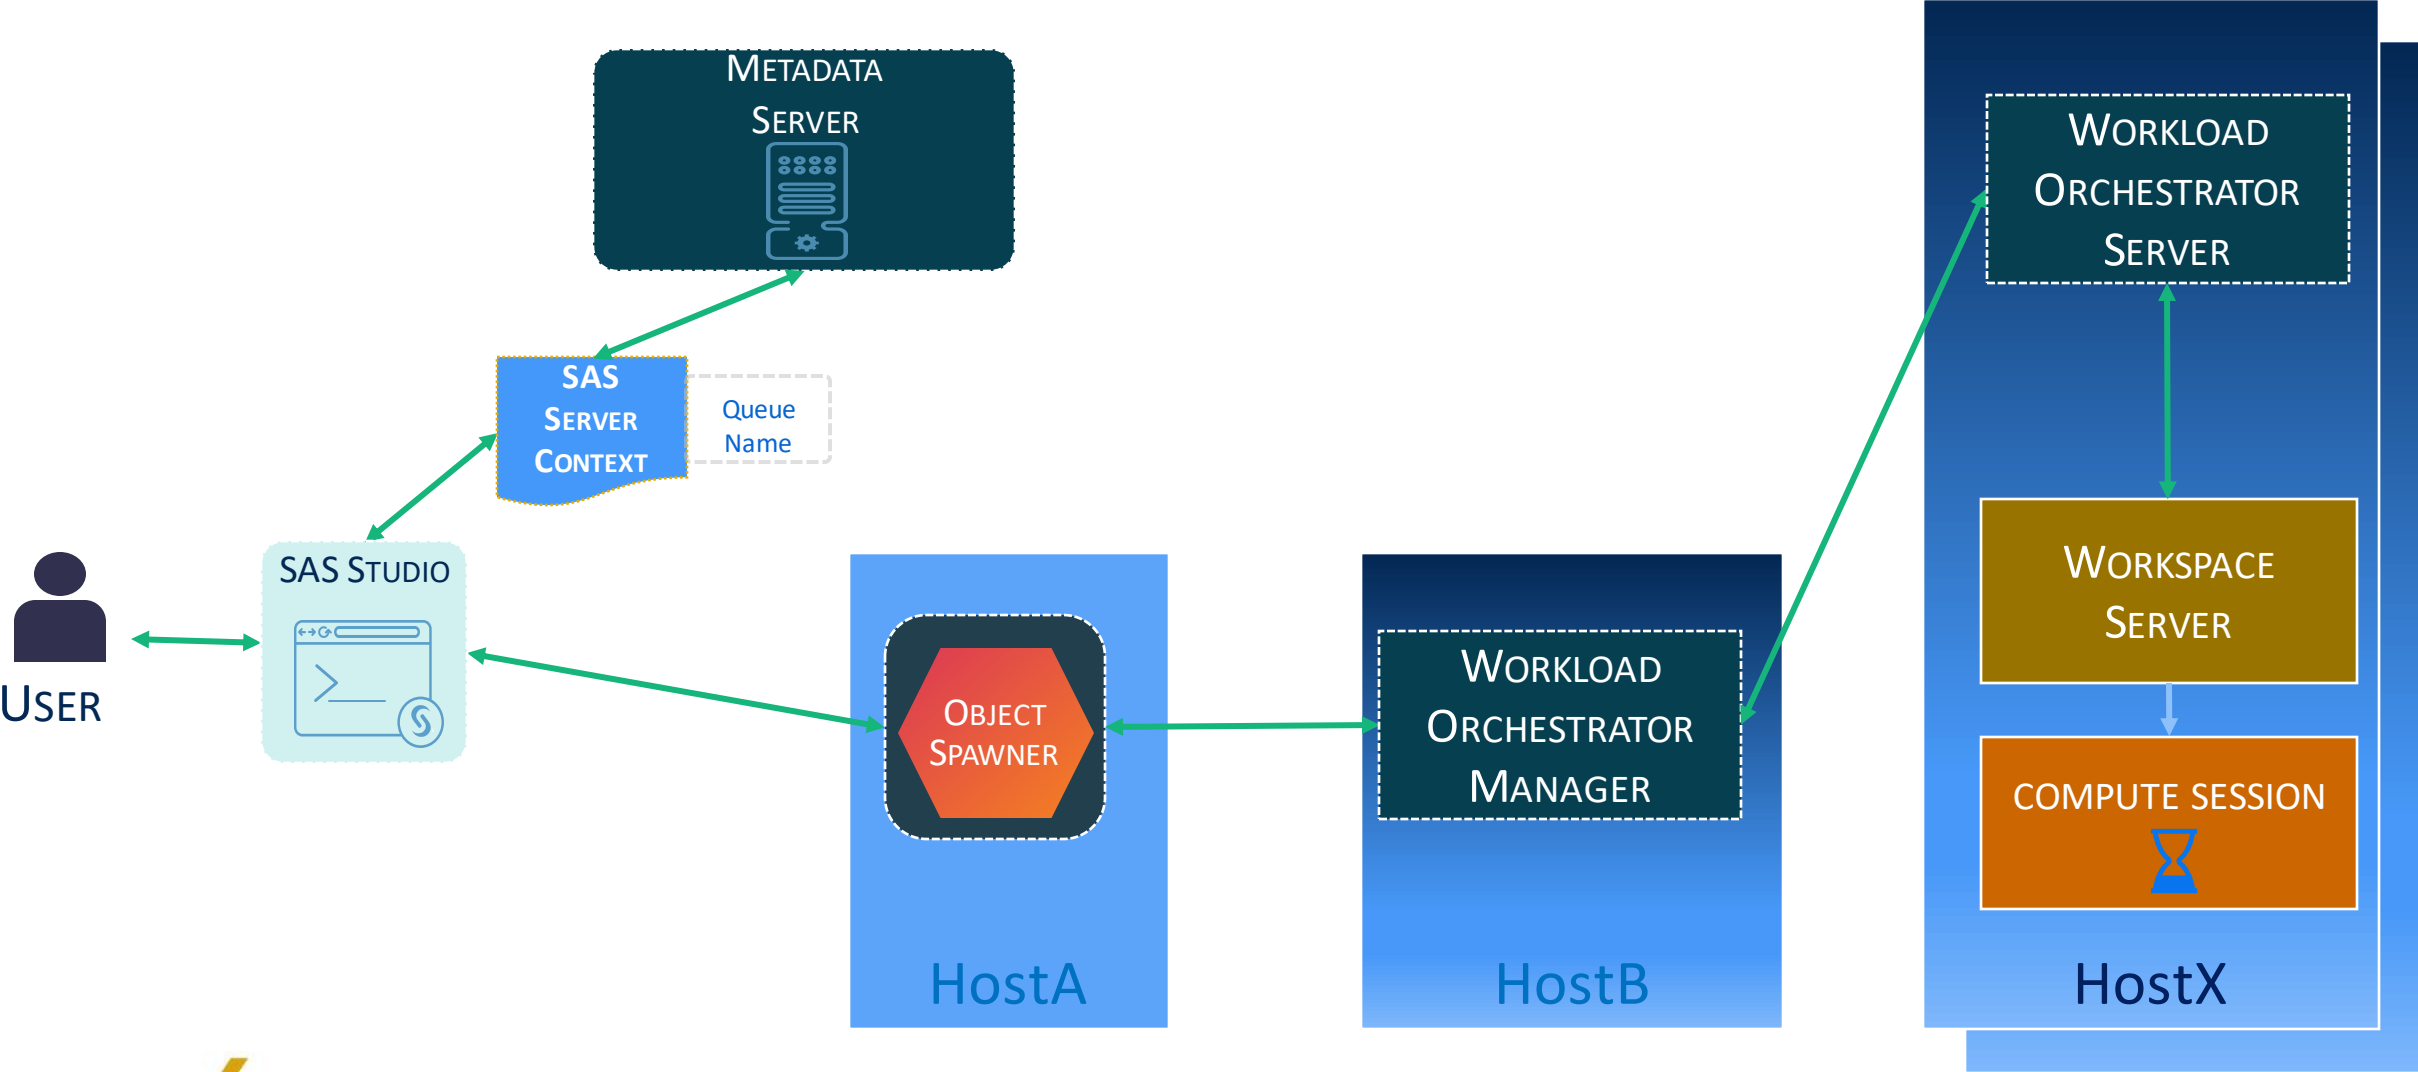


A Brief History of Workload Management...

SAS 9 LOAD BALANCED WORKSPACE SERVERS



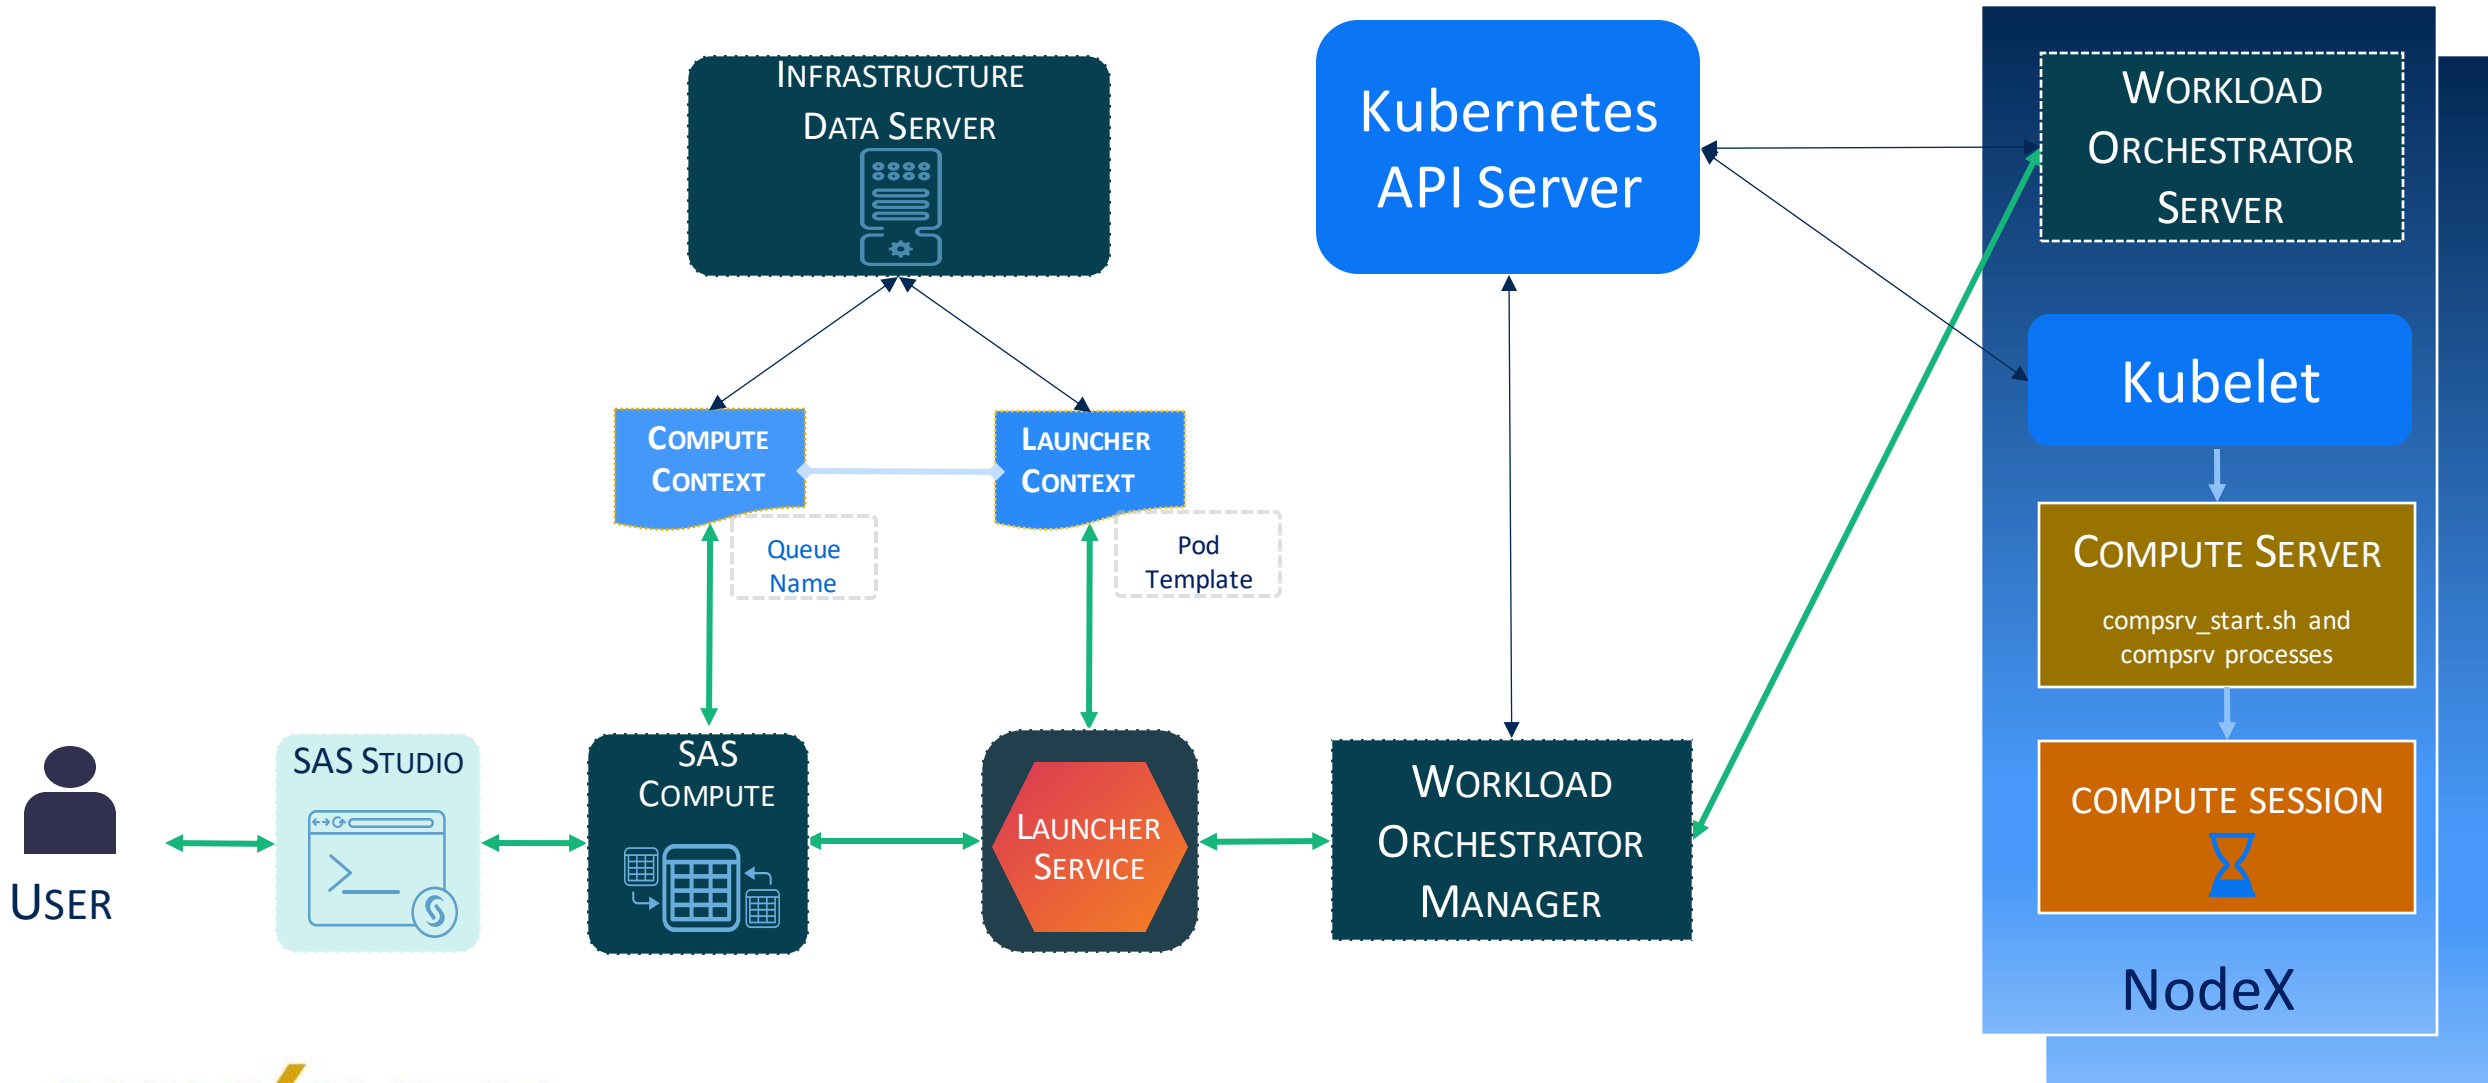
SAS 9.4 GRID-LAUNCHED WORKSPACE SERVERS





Workload Management in SAS Viya...

SAS VIYA 4 SAS WORKLOAD MANAGEMENT



Demo of Workload Management in SAS Viya

Autoscaling Configuration Summary

#1 - Node Pool Definition

Node pool	
Max pods per node	110
Public IPs per node	Disabled
Autoscaling	Enabled
Azure Spot Instance	Disabled

Taints and labels	
Taints	<code>workload.sas.com/class=compute:NoSchedule</code>
Labels	<code>demo.sas.com/type : gpu</code> <code>launcher.sas.com/prepullimage : sas-programming-environment</code> <code>workload.sas.com/class : compute</code>

#2 - Host Type Definition

Host type name:	<code>gpuHostType</code>
Host properties:	<code>demo.sas.com/type=gpu</code>
<input checked="" type="checkbox"/> Enable Autoscaling	

#4 - Context Definitions

Name:	SAS Studio launcher context
Name:	SAS Studio compute context
Name:	default
Description:	Default Batch Context

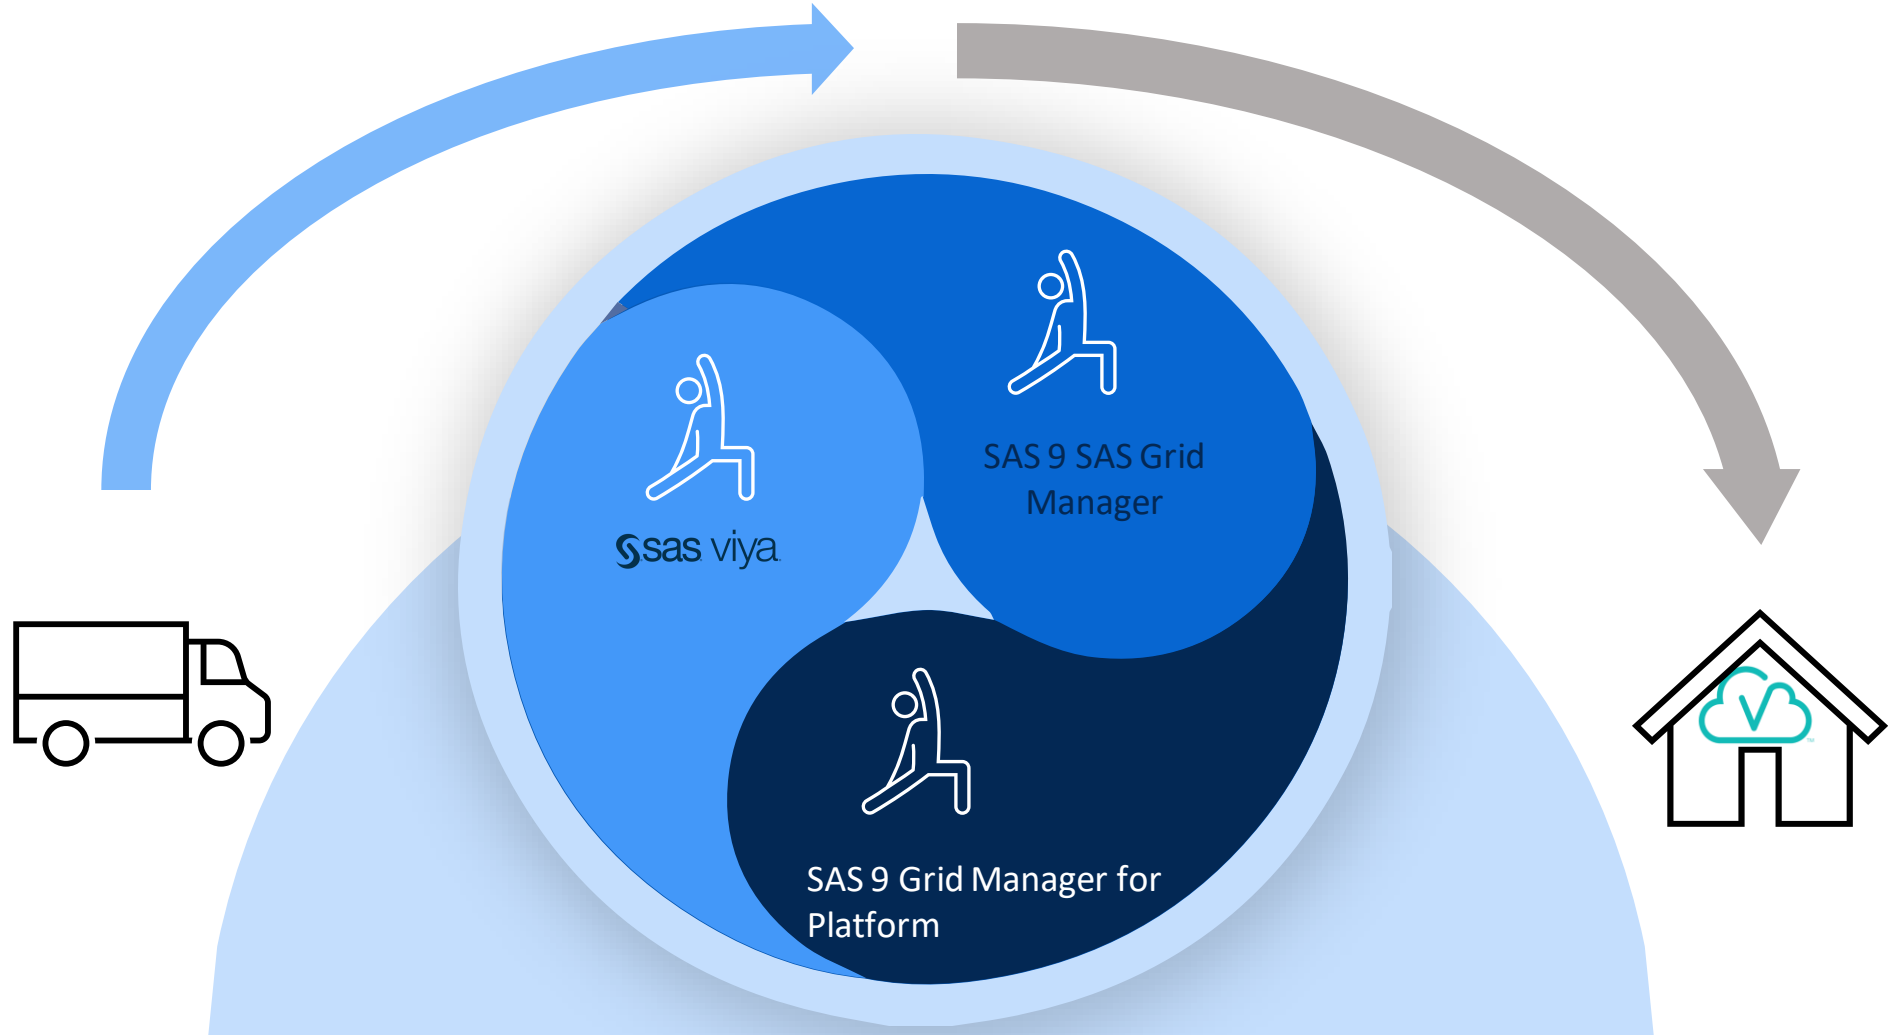
SAS Workload Orchestrator queue: `gpuQueue`

#3 - Queue Definition

Name:	<code>gpuQueue</code>
Host types:	<code>gpuHostType</code>
Autoscaler minimum pending jobs:	5
Autoscaler minimum pending time for jobs (seconds):	60

SAS Workload Management


Grid Migration



AGENDA



Motivation



Capabilities



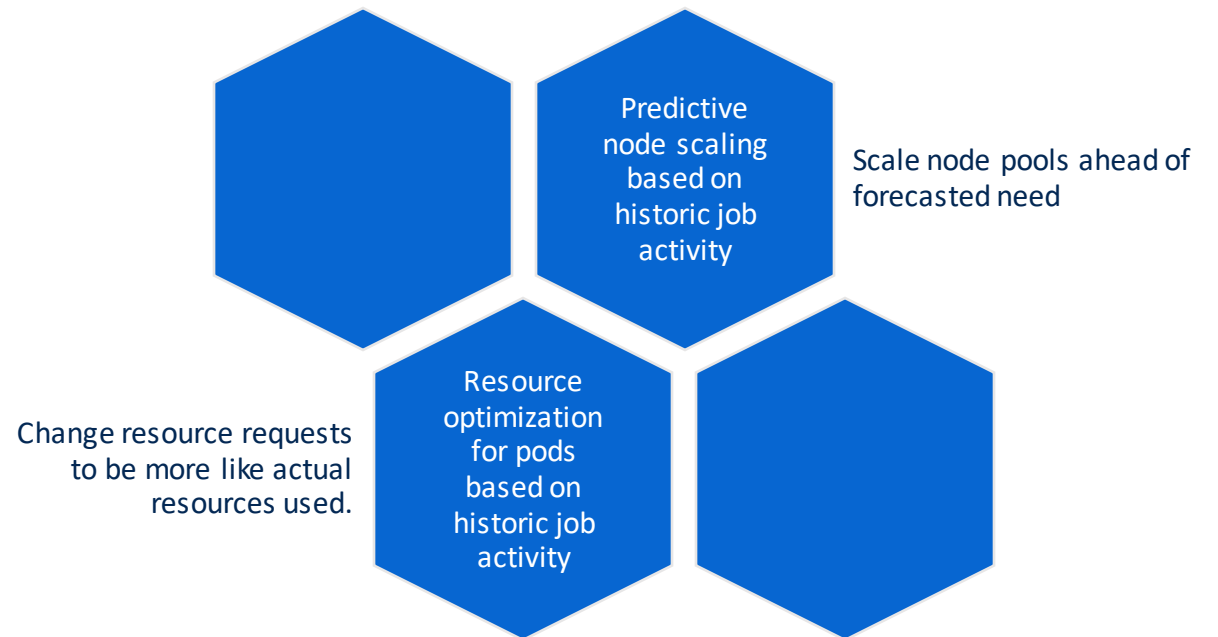
Demo



Future




Analytics






Visualization

- Update UI with visuals related to
 - Host Information
 - % Utilization
 - Job Information
 - Queue Information
 - Running jobs
 - Pending jobs
 - Job information
 - Gantt chart of running jobs start time (end time) relative to current time
- 




Thank you!

Alex.Daehnrich@sas.com
Doug.Haigh@sas.com



[Try SAS Viya for Free](https://www.sas.com/en_us/trials/software/viya/viya-trial-form.html)

https://www.sas.com/en_us/trials/software/viya/viya-trial-form.html



SAS **EXPLORE**


Copyright © SAS Institute Inc. All rights reserved.

sas

Backup Slides for SWO UI

Contexts

View: Compute contexts ▾



- CAS Formats service compute context
- CAS Management service compute context
- Import 9 service compute context
- SAS Backup job compute context
- SAS Job Execution compute context
- SAS Studio compute context
- SAS Studio compute context for Priority Users**

Basic Properties

> Resources

▾ Advanced

Created by:

dohaig

Date created:

July 27, 2023 01:12:08 PM

Modified by:

dohaig

Date modified:

July 27, 2023 01:12:08 PM

Environment

SAS Workload Orchestrator queue:


InteractivePriority

SAS options:

Contexts

View:

Compute contexts ▾



- CAS Formats service compute context
- CAS Management service compute context
- Import 9 service compute context
- SAS Backup job compute context
- SAS Job Execution compute context
- SAS Studio compute context
- SAS Studio compute context for Priority Users**

Basic Properties

Name:

SAS Studio compute context for Priority Users

Description:

Compute context to be used by the SAS Studio service for Priority Users

Launch type:

Service

Launcher context:

SAS Studio launcher context for Priority Users

Identities:


Authenticated users

Attributes

Attribute	Value
-----------	-------

Contexts

View: Launcher contexts ▾



- CAS Formats service launcher context
- CAS Management launcher context
- Import 9 service launcher context
- QKB Bootstrap launcher context
- SAS Admin Content Loader launcher context
- SAS Backup job launcher context
- SAS Batch service launcher context
- SAS Batch service launcher context for commands
- SAS Batch service launcher context for Priority Users
- SAS Job Execution launcher context
- SAS Studio launcher context
- SAS Studio launcher context for Priority Users**
- SAS/CONNECT service launcher context

Basic Properties

Name:

SAS Studio launcher context for Priority Users

Description:

Launcher context to be used by the SAS Studio service for Priority Users

Pod template name:

sas-compute-job-config

Environment Variables

Variable	Value
SAS_LAUNCHER_INIT_ENCODING_DEFAULT	utf8
SAS_LAUNCHER_INIT_LOCALE_DEFAULT	en_US

> Commands

> Advanced

Workload Orchestrator

Dashboard Jobs Queues Hosts Logs Configuration Log Levels

Information

General

Version:

1.19.6

License expiration date:

Oct/25/2023

Build date:

July 25, 2023 06:08:38 PM

GUI build date (version):

July 25, 2023 04:55:04 AM
(5.11.28)

Jobs

Total:

Queue Status



default

Priority 10

OPEN-ACTIVE

Jobs pending: 0

Jobs running: 0

Host Status



cnt-kc1-n1

9% Utilized



cnt-kc1-n2

4% Utilized

Workload Orchestrator

Dashboard **Jobs** Queues Hosts Logs Configuration Log Levels

Filter



Jobs (80)

<input type="checkbox"/>	ID	State	User	Name	Queue	Tenant	Execu...	Subm...	Start ...	End T...
<input type="checkbox"/>	904	RUNNING	dohaig	DougTe st_7	default	uaa	cnt-kc1- n3	2023 11:53:0 5 AM	2023 11:53:0 6 AM	-
<input checked="" type="checkbox"/>	905	RUNNING	dohaig	DougTe st_13	default	uaa	cnt-kc1- n1	July 27, 2023 11:53:0 6 AM	July 27, 2023 11:53:0 6 AM	-
<input type="checkbox"/>	906	RUNNING	dohaig	DougTe st_8	default	uaa	cnt-kc1- n1	July 27, 2023 11:53:0 6 AM	July 27, 2023 11:53:0 6 AM	-
<input type="checkbox"/>	907	RUNNING	dohaig	DougTe st_21	default	uaa	cnt-kc1- n1	July 27, 2023 11:53:0 7 AM	July 27, 2023 11:53:0 7 AM	-

View: 50 ◾ ◀ Displaying 1-50 ▶

Workload Orchestrator

Dashboard Jobs Queues Hosts Logs Configuration Log Levels

Jobs > **DougTest_13**



Current Status

Host Information

Information

Limits

Output

Statistics

ID:	905
Name:	DougTest_13
User:	dohaig
Queue:	default
State:	RUNNING
Pending cause:	
Submit time:	July 27, 2023 11:53:06 AM
Start time:	July 27, 2023 11:53:06 AM
End time:	-

Workload Orchestrator

Dashboard Jobs Queues Hosts Logs Configuration Log Levels

Jobs > DougTest_13



Current Status

Host Information

Information

Limits

Output

Statistics

Filter

Type	Name	Value
Final consumed resources	memory	536.102295
Host information	Execution host	cnt-kc1-n1
Host information	Execution Host Alias	sas-workload-orchestrator-c7tgq
Host information	Exit code	0
Host information	Launch status	0
Host information	Process ID (PID)	118
Host scheduling information	cnt-kc1-n1	The host cnt-kc1-n1 is available to run the job.
Host scheduling information	cnt-kc1-n2	The host cnt-kc1-n2 is available to run the job.
Host scheduling information	cnt-kc1-n3	The host cnt-kc1-n3 is available to run the job.
		The host cnt-kc1-n4 is available to

Workload Orchestrator

Dashboard Jobs Queues Hosts Logs Configuration Log Levels

Filter

🔒 | Activate | Inactivate | ↻

Default queue: default

<input type="checkbox"/>	Name	Tenant	State	Priority	Jobs Pen...	Jobs Run...	Maximum...	Maximum... ⚑
<input checked="" type="checkbox"/>	<u>default</u>	uaa	OPEN- ACTIVE	10	0	36	Unde...	Unde...
<input type="checkbox"/>	<u>BatchNor...</u>	uaa	OPEN- ACTIVE	20	0	0	4	Unde...
<input type="checkbox"/>	<u>BatchPriority</u>	uaa	OPEN- ACTIVE	30	0	0	Unde...	Unde...
<input type="checkbox"/>	<u>PolicyDe...</u>	uaa	OPEN- ACTIVE	5	0	0	15	10
<input type="checkbox"/>	<u>Interactiv...</u>	uaa	OPEN- ACTIVE	40	0	0	Unde...	Unde...
<input type="checkbox"/>	<u>Interactiv...</u>	uaa	OPEN- ACTIVE	50	0	0	Unde...	Unde...

Workload Orchestrator

Dashboard Jobs Queues Hosts Logs Configuration Log Levels

Queues > default

🔒 🔒 Activate Inactivate ↻

General

Name:

default

Tenant:

uaa

Priority:

10

State:

OPEN-ACTIVE

Default queue:

true

Jobs in default queue

🔍 Filter

⏸ ⏪ ⏩ ↻

<input type="checkbox"/>	ID	State	User	Name	Execu...	Subm...	Start Time	End Ti...	⌵
<input type="checkbox"/>	858	RUNNING	dohaig	DougTe st_2	cnt-kc1- n2	July 27, 2023 11:47:3 8 AM	July 27, 2023 11:47:38 AM	-	
<input type="checkbox"/>	859	RUNNING	dohaig	DougTe st_21	cnt-kc1- n2	July 27, 2023 11:47:3 8 AM	July 27, 2023 11:47:38 AM	-	
<input type="checkbox"/>	860	RUNNING	dohaig	DougTe st_5	cnt-kc1- n2	July 27, 2023 11:47:3 8 AM	July 27, 2023 11:47:38 AM	-	
<input type="checkbox"/>	861	RUNNING	dohaig	DougTe st_9	cnt-kc1- n2	July 27, 2023 11:47:3 9 AM	July 27, 2023 11:47:39 AM	-	

Workload Orchestrator

Dashboard Jobs Queues **Hosts** Logs Configuration Log Levels

Hosts > cnt-kc1-n1



General

CPU

Jobs (Overview)

Jobs (Running)

Memory

OS

Resources

Schedule Thresholds

Suspend Thresholds

Host name:	cnt-kc1-n1
Host type:	default
State:	OPEN-FULL
Build date:	July 25, 2023 06:08:38 PM
Operating system name:	Linux
Operating system description:	Red Hat Enterprise Linux 8
CPU vendor:	Intel
CPU architecture:	x86_64
Cores:	12

Workload Orchestrator

Dashboard Jobs Queues **Hosts** Logs Configuration Log Levels

Hosts > cnt-kc1-n1



General

CPU

Jobs (Overview)

Jobs (Running)

Memory

OS

Resources

Schedule Threshold

Suspend Threshold

Jobs on host cnt-kc1-n1

Filter



<input type="checkbox"/>	ID	State	User	Name	Queue	Submit ...	Start Time	End Time	
<input type="checkbox"/>	874	STARTING	dohaig	DougTest _36	default	July 27, 2023 11:47:44 AM	July 27, 2023 11:47:44 AM	-	
<input type="checkbox"/>	875	STARTING	dohaig	DougTest _7	default	July 27, 2023 11:47:44 AM	July 27, 2023 11:47:44 AM	-	
<input type="checkbox"/>	876	RUNNING	dohaig	DougTest _16	default	July 27, 2023 11:47:44 AM	July 27, 2023 11:47:44 AM	-	
<input type="checkbox"/>	877	STARTING	dohaig	DougTest _22	default	July 27, 2023 11:47:44 AM	July 27, 2023 11:47:44 AM	-	

Workload Orchestrator

Dashboard Jobs Queues Hosts **Logs** Configuration Log Levels

Filter by message

Advanced Filter

▼ Logger (no filter)

- App.Grid.SGMG
- App.Grid.SGMG....
- App.Grid.SGMG.L...
- App.Grid.SGMG.L...
- App.Grid.SGMG.L...
- App.Grid.SGMG.L...

▼ Level (1 of 6)

- FATAL
- ERROR
- WARN

Apply

```
2023-07-27T15:08:27Z INFO [00000005] sas -
2023-07-27T15:08:27Z INFO [00000005] sas - SAS Workload Orchestrator Daemon - Version 4.0 (2023-07
2023-07-27T15:08:27Z INFO [00000005] sas - Copyright (C) 2017-2023, SAS Institute Inc., Cary, NC, USA.
2023-07-27T15:08:27Z INFO [00000005] sas - SAS Workload Orchestrator running in Kubernetes with cor
2023-07-27T15:08:27Z INFO [00000005] sas -
2023-07-27T15:08:27Z INFO [00000008] sas - SGMG is running under the user identity .
2023-07-27T15:08:27Z INFO [00000008] sas - SGMG has completed initialization.
2023-07-27T15:08:28Z INFO [00000008] sas - The host "sas-workload-orchestrator-0" has been added to
2023-07-27T15:08:28Z INFO [00000008] sas - The host "sas-workload-orchestrator-1" has been added to
2023-07-27T15:08:28Z INFO [00000008] sas - The host "sas-workload-orchestrator-2" has been added to
2023-07-27T15:08:28Z INFO [00000008] sas - The host "sas-workload-orchestrator-3" has been added to
2023-07-27T15:08:28Z INFO [00000008] sas - The host "sas-workload-orchestrator-4" has been added to
2023-07-27T15:08:28Z INFO [00000008] sas - The host "sas-workload-orchestrator-5" has been added to
2023-07-27T15:08:28Z INFO [00000008] sas - The host "sas-workload-orchestrator-6" has been added to
2023-07-27T15:08:28Z INFO [00000008] sas - The host "sas-workload-orchestrator-7" has been added to
2023-07-27T15:08:28Z INFO [00000008] sas - The host "sas-workload-orchestrator-8" has been added to
2023-07-27T15:08:28Z INFO [00000008] sas - The host "sas-workload-orchestrator-9" has been added to
2023-07-27T15:08:28Z INFO [00000008] sas - The queue "default" has been added to the scheduler.
2023-07-27T15:08:28Z INFO [00000008] sas - Starting HTTP server 'sas-workload-orchestrator-0', minPort
2023-07-27T15:08:28Z INFO [00000008] sas - HTTP Server 'sas-workload-orchestrator-0' listening on port
2023-07-27T15:08:29Z ERROR [00000011] sas - ERROR: Error connecting to 192.168.247.82:8080. (The c
2023-07-27T15:08:29Z ERROR [00000011] sas - ERROR: Unable to establish connection to 192.168.247.8

```

Workload Orchestrator

Dashboard Jobs Queues Hosts Logs Configuration **Log Levels**

Filter

Set log level



<input type="checkbox"/>	Logger Name ↑	Level	Level Description	
<input type="checkbox"/>	Admin	NULL	Inherit parent logging level	
<input type="checkbox"/>	Admin.Operations	NULL	Inherit parent logging level	
<input type="checkbox"/>	App	NULL	Inherit parent logging level	
<input type="checkbox"/>	App.Grid	NULL	Inherit parent logging level	
<input type="checkbox"/>	App.Grid.SGMG	NULL	Inherit parent logging level	
<input type="checkbox"/>	App.Grid.SGMG.Audit	INFO	Log FATAL, ERROR, WARN, and INFO messages	
<input type="checkbox"/>	App.Grid.SGMG.Log	NULL	Inherit parent logging level	
<input type="checkbox"/>	App.Grid.SGMG.Log.Config	NULL	Inherit parent logging level	
<input type="checkbox"/>	App.Grid.SGMG.Log.Config.Date	NULL	Inherit parent logging level	
<input type="checkbox"/>	App.Grid.SGMG.Log.Config.JSON	NULL	Inherit parent logging level	
<input type="checkbox"/>	App.Grid.SGMG.Log.DB	NULL	Inherit parent logging level	

Workload Orchestrator

Dashboard Jobs Queues Hosts Logs Configuration Log Levels

General

Compare Order

Host Types

Queues

User Groups

User Resources

Filter



New queue



Name: PolicyDemoQueue

Tenant: uaa

Description: Queue to show off various workload management policy feature:

Default Time Based Settings Night +

Priority: * 5

Default queue

Restart jobs

Maximum jobs: 15

Maximum jobs per user: 5

Workload Orchestrator

Dashboard Jobs Queues Hosts Logs Configuration Log Levels

General

Compare Order

Host Types

Queues

User Groups

User Resources

Filter



New queue



Name: PolicyDemoQueue

Tenant: uaa

Description: Queue to show off various workload management policy feature:

Default Time Based Settings Night +

Priority: * 5

Default queue

Restart jobs

Maximum jobs: 15

Maximum jobs per user: 5

Workload Orchestrator

Dashboard Jobs Queues Hosts Logs Configuration Log Levels

General

Compare Order

Host Types

Queues

User Groups

User Resources

Filter

New queue

Users:

sascnn4 x

sascnn3 x

sascnn2 x

sascnn1 x

sascnn x

Administrators:

normal-q-admins x

Host names:

Host types:

PolicyDemoHostType x

Required tags:

GPU x

Preempts:

Autoscaler minimum pending jobs:

Autoscaler minimum pending time for jobs (seconds):

Workload Orchestrator

Dashboard Jobs Queues Hosts Logs Configuration Log Levels

General

Compare Order

Host Types

Queues

User Groups

User Resources

Filter

New queue

Name: PolicyDemoQueue

Tenant: uaa

Description: Queue to show off various workload management policy features

Default Time Based Settings Night +

Priority: * 5

Default queue

Restart jobs

Maximum jobs: 15

Maximum jobs per user: 5

Maximum jobs per host: 20

Workload Orchestrator

Dashboard Jobs Queues Hosts Logs Configuration Log Levels

- General
- Compare Order
- Host Types
- Queues
- User Groups
- User Resources

Filter

New queue

CONSUMED RESOURCES

cores:

memory:

Select an item +

REQUIRED RESOURCES

Select an item +

LIMITS

maxClockTime:

Select an item +

Workload Orchestrator

Dashboard Jobs Queues Hosts Logs Configuration Log Levels

General

Compare Order

Host Types

Queues

User Groups

User Resources

Filter

New queue

Users:

sascnn4 x

sascnn3 x

sascnn2 x

sascnn1 x



sascnn x

Administrators:

normal-q-admins x



Host names:



Host types:

PolicyDemoHostType x



Required tags:

GPU x



Preempts:



Autoscaler minimum pending jobs:

Autoscaler minimum pending time for jobs (seconds):

Workload Orchestrator

Dashboard Jobs Queues Hosts Logs Configuration Log Levels

General

Compare Order

Host Types

Queues

User Groups

User Resources

Filter

New queue

Users:

sascnn4 x

sascnn3 x

sascnn2 x

sascnn1 x

sascnn x

Administrators:

normal-q-admins x

Host names:

Host types:

PolicyDemoHostType x

Required tags:

GPU x

Preempts:

Autoscaler minimum pending jobs:

Autoscaler minimum pending time for jobs (seconds):

Workload Orchestrator

Dashboard Jobs Queues Hosts Logs Configuration Log Levels

- General
- Compare Order
- Host Types
- Queues
- User Groups
- User Resources

Filter

New queue

CONSUMED RESOURCES

cores:



memory:



Select an item +

REQUIRED RESOURCES

Select an item +

LIMITS

maxClockTime:



Select an item +

Workload Orchestrator

Dashboard Jobs Queues Hosts Logs Configuration Log Levels

General

Compare Order

Host Types

Queues

User Groups

User Resources

Filter

New queue

Users:

sascnn4 x

sascnn3 x

sascnn2 x

sascnn1 x

sascnn x

Administrators:

normal-q-admins x

Host names:

Host types:

PolicyDemoHostType x

Required tags:

GPU x

Preempts:

Autoscaler minimum pending jobs:

Autoscaler minimum pending time for jobs (seconds):

Workload Orchestrator

Dashboard Jobs Queues Hosts Logs Configuration Log Levels

General

Compare Order

Host Types

Queues

User Groups

User Resources

Filter

New queue

Default Time Based Settings

Night

Priority:

Default queue

Default queue:

Use default

Restart jobs:

Use default

Maximum jobs:

Maximum jobs per user:

Maximum jobs per host:

20

Enable users override

Workload Orchestrator

Dashboard Jobs Queues Hosts Logs Configuration Log Levels

General

Compare Order

Host Types

Queues

User Groups

User Resources

Filter

PolicyDemoHostType

Host type name:

PolicyDemoHostType

Description:

Host Type to demonstrate workload management host types

Host names:

Host name patterns:


Host properties:

mylabel.sas.com=policydemo x


IP patterns:

IP ranges:

Enable Autoscaling



New host type



Workload Orchestrator

Dashboard Jobs Queues Hosts Logs Configuration Log Levels

General

Compare Order

Host Types

Queues

User Groups

User Resources

Filter

PolicyDemoHostType

Host type name:

PolicyDemoHostType

Description:

Host Type to demonstrate workload management host types

Host names:

Host name patterns:

Host properties:

mylabel.sas.com=policydemo x

IP patterns:

IP ranges:

Enable Autoscaling



New host type



Workload Orchestrator

Dashboard Jobs Queues Hosts Logs Configuration Log Levels

General

Compare Order

Host Types

Queues

User Groups

User Resources

Filter

PolicyDemoHostType

Host type name:

PolicyDemoHostType

Description:

Host Type to demonstrate workload management host types

Host names:

Host name patterns:

Host properties:


mylabel.sas.com=policydemo x

IP patterns:

IP ranges:

Enable Autoscaling

New host type



Workload Orchestrator

Dashboard Jobs Queues Hosts Logs Configuration Log Levels



- General
- Compare Order
- Host Types
- Queues
- User Groups
- User Resources

Filter

New host type

Default Time Based Settings +

Maximum jobs allowed: * 20

Tags:

SCHEDULE THRESHOLDS

utilization: > 0.85 <= 0.7

Select an item +

SUSPEND THRESHOLDS

utilization: > 0.95 <= 0.85

Select an item +

Backup Slides for Auto-Scaling Configuration






Home > dohaig-aks-aks | Node pools >

compgpu | Overview

Node pool



 Start Stop  Upgrade Kubernetes  Update

 Overview

 Nodes

 Configuration

Node pool

Max pods per node	110
Public IPs per node	Disabled
Autoscaling	Disabled
Azure Spot Instance	Disabled
Maximum price	N/A
Scale eviction policy	N/A
Node image version	AKSUbuntu-2204gen2containerd-202308.01.0
Proximity placement group	N/A

Home > dohaig-aks-aks | Node pools >

compgpu | Overview

Node pool

Search

Start Stop Upgrade Kubernetes

- Overview
- Nodes
- Configuration

Essentials

Provisioning state ⓘ
Succeeded

Power state ⓘ
Running (0/0 nodes ready)

Availability zones
None

Mode
User

Scale node pool

compgpu

You can scale the number of nodes in your cluster to increase the total amount of cores and memory available for your container applications. [Learn more](#)

Scale method ⓘ

Manual

Autoscale - **Recommended**

This option is recommended so that the cluster is automatically sized correctly for the current running workloads.

Minimum node count * ⓘ

Maximum node count * ⓘ

The maximum node count allowed for an AKS cluster is 1000 per node pool and 5000 nodes across all node pools in this cluster.






Home > dohaig-aks-aks | Node pools >

compgpu | Overview ...


Node pool



 Start Stop  Upgrade Kubernetes  Update

 Overview

 Nodes

 Configuration

Node pool

Max pods per node	110
Public IPs per node	Disabled
Autoscaling	Enabled
Azure Spot Instance	Disabled
Maximum price	N/A
Scale eviction policy	N/A
Node image version	AKSUbuntu-2204gen2containerd-202308.01.0
Proximity placement group	N/A




Home > dohaig-aks-aks | Node pools >


compgpu | Overview

Node pool

[▶ Start](#) [□ Stop](#) [↑ Upgrade Kubernetes](#) [↑ Update image](#) [↗ Scale node](#)

 Overview

 Nodes

 Configuration

Taints and labels

Taints

`workload.sas.com/class=compute:NoSchedule`

Labels

`demo.sas.com/type : gpu`

`launcher.sas.com/prepareimage : sas-programming-environment`

`workload.sas.com/class : compute`

Workload Orchestrator

Dashboard Jobs Queues Hosts Logs Configuration Log Levels

General

Compare Order

Host Types

Queues

User Groups

User Resources

Filter

ScalableHostTypeGPU

Host type name:

ScalableHostTypeGPU

Description:

Host names:

Host name patterns:

Host properties:

demo.sas.com/type=gpu x

IP patterns:

IP ranges:




Enable Autoscaling

New host type

Workload Orchestrator

Dashboard Jobs Queues Hosts Logs Configuration Log Levels



New queue  


General

Compare Order

Host Types

Queues

User Groups


User Resources

▼ scalable-gpu


Name: gpu

Tenant: uaa

Description:

 Default Time Based Settings +

Priority: *


Host types: 

Autoscaler minimum pending jobs:

Autoscaler minimum pending time for jobs (seconds):

Contexts

View: Compute contexts ▾



- CAS Formats service compute context
- CAS Management service compute context
- Import 9 service compute context
- SAS Backup job compute context
- SAS Job Execution compute context
- SAS Studio compute context**
- SAS Studio compute context for Priority Users

Basic Properties

> Resources

▾ Advanced

Created by:

dohaig

Date created:

July 27, 2023 01:12:08 PM

Modified by:

dohaig

Date modified:

July 27, 2023 01:12:08 PM

Environment

SAS Workload Orchestrator queue:

gpu

SAS options:

Backup Slides for Auto-Scaling

SAS WORKLOAD ORCHESTRATOR CLUSTER AUTOSCALING

Demonstration of Kubernetes Activity

At start – no SWO daemons

```
sas-batch-6bc9f9d76c-ncmth    1/1 Running aks-stateless-41936518-vmss00000n
sas-workload-orchestrator-0  1/1 Running aks-stateful-10864146-vmss00000n
sas-workload-orchestrator-1  1/1 Running aks-stateless-41936518-vmss00000n
```

Jobs submitted; SWO attempts to cause scaling by creating scaling pod

```
sas-batch-6bc9f9d76c-ncmth    1/1 Running aks-stateless-41936518-vmss00000n
sas-workload-orchestrator-0  1/1 Running aks-stateful-10864146-vmss00000n
sas-workload-orchestrator-1  1/1 Running aks-stateless-41936518-vmss00000n
sas-workload-orchestrator-scaling-pod-00007f102a7b8ea0 0/1 Pending <none>
```

SAS WORKLOAD ORCHESTRATOR CLUSTER AUTOSCALING

Demonstration of Kubernetes Activity

Node created; sas-prepull starts

```

sas-batch-6bc9f9d76c-ncmth      1/1 Running      aks-stateless-41936518-vmss00000n
sas-prepull-image-pod-ulsijuvw  0/1 ContainerCreating aks-compute2-15904406-vmss00007n
sas-workload-orchestrator-0     1/1 Running      aks-stateful-10864146-vmss00000n
sas-workload-orchestrator-1     1/1 Running      aks-stateless-41936518-vmss00000n
sas-workload-orchestrator-scaling-pod-00007f102a7b8ea0 0/1 Pending      <none>

```

SWO daemonset pod created on node and scaling pod assigned to node

```

sas-batch-6bc9f9d76c-ncmth      1/1 Running      aks-stateless-41936518-vmss00000n
sas-prepull-image-pod-ulsijuvw  0/1 ContainerCreating aks-compbigmem-15904406-vmss00007n
sas-workload-orchestrator-0     1/1 Running      aks-stateful-10864146-vmss00000n
sas-workload-orchestrator-1     1/1 Running      aks-stateless-41936518-vmss00000n
sas-workload-orchestrator-pwnk7  0/1 Init:0/2     aks-compbigmem-15904406-vmss00007n
sas-workload-orchestrator-scaling-pod-00007f102a7b8ea0 0/1 Init:0/1     aks-compbigmem-15904406-vmss00007n

```


SAS WORKLOAD ORCHESTRATOR CLUSTER AUTOSCALING

Demonstration of Kubernetes Activity

SWO daemonset pod finishes initializing; scaling pod deleted; prepull completed

```

sas-batch-6bc9f9d76c-ncmth      1/1 Running aks-stateless-41936518-vmss00000n
sas-workload-orchestrator-0    1/1 Running aks-stateful-10864146-vmss00000n
sas-workload-orchestrator-1    1/1 Running aks-stateless-41936518-vmss00000n
sas-workload-orchestrator-pwnk7 0/1 Running aks-compute2-15904406-vmss00007n

```

SWO manager schedules jobs to new node

```

sas-batch-6bc9f9d76c-ncmth      1/1 Running aks-stateless-41936518-vmss00000n
sas-batch-server-0c0d8c71-2c63-4cca-8965-98f4e2dcc56b-86 0/2 Init:0/2 aks-compbigmem-15904406-vmss00007n
sas-batch-server-0c1b662b-3f54-47ab-8cd8-5fd2dddde14a-85 0/2 Init:0/2 aks-compbigmem-15904406-vmss00007n
sas-batch-server-151e8db6-69a3-41d9-80d5-469fbf413ac1-87 0/2 Init:0/2 aks-compbigmem-15904406-vmss00007n
sas-batch-server-3161a16c-8e0e-4b7b-8960-ca52c5b24646-91 0/2 Pending aks-compbigmem-15904406-vmss00007n
sas-batch-server-97ccea36-fcb8-46fb-94bc-ea6fb95f70d8-89 0/2 Pending aks-compbigmem-15904406-vmss00007n
sas-batch-server-a68f40ca-4a4f-43b9-a7e7-f4886b50f13a-90 0/2 Init:0/2 aks-compbigmem-15904406-vmss00007n
sas-batch-server-d3a0c019-7bc7-4a2a-98a5-9cadf169a564-92 0/2 Pending aks-compbigmem-15904406-vmss00007n
sas-batch-server-d74dedaa-1881-4836-95b4-b2144875abd5-88 0/2 Pending aks-compbigmem-15904406-vmss00007n
sas-workload-orchestrator-0    1/1 Running aks-stateful-10864146-vmss00000n
sas-workload-orchestrator-1    1/1 Running aks-stateless-41936518-vmss00000n
sas-workload-orchestrator-pwnk7 0/1 Running aks-compbigmem-15904406-vmss00007n

```

Backup Slides for Grid Configuration Migration

SASv9 GRID CONFIGURATION MIGRATION

Convert/Import LSF Configuration From LSF Configuration Directory

```
$> ./sas-viya workload-orchestrator config grid convert-from-lsf-config-dir \  
--cluster-name test_cluster \  
--config-path ./lsf-config \  
> lsf-config.json
```

The conversion of the LSF configuration to the SAS Workload Orchestrator configuration has resulted in the following messages:

```
* Warning: SAS Workload Orchestrator does not support time-based configurations for user groups.  
* Warning: SAS Workload Orchestrator does not support time-based configurations for host groups.  
* Warning: SAS Workload Orchestrator does not support the complex resource requirement "select[ poe > 0 ]" that  
is found in the "hpc_ibm" queue configuration. The setting is ignored.  
* Warning: SAS Workload Orchestrator does not support the complex resource requirement "select[ poe > 0 ]" that  
is found in the "hpc_ibm_tv" queue configuration. The setting is ignored.
```

SASv9 GRID CONFIGURATION MIGRATION

Convert/Import LSF Configuration From Platform Web Services

```
$> ./sas-viya workload-orchestrator config grid convert-from-pws \  
      --config-id 243 \  
      --url http://myPwsHost.mycompany.com \  
      --user lsfGridAdmin \  
      --pwd lsfGridAdminPwd \  
> lsf-config.json
```

The conversion of the LSF configuration to the SAS Workload Orchestrator configuration has resulted in the following messages:

```
* Warning: SAS Workload Orchestrator does not support time-based configurations for user groups.  
* Warning: SAS Workload Orchestrator does not support time-based configurations for host groups.  
* Warning: SAS Workload Orchestrator does not support the complex resource requirement "select[ poe > 0 ]" that  
is found in the "hpc_ibm" queue configuration. The setting is ignored.  
* Warning: SAS Workload Orchestrator does not support the complex resource requirement "select[ poe > 0 ]" that  
is found in the "hpc_ibm_tv" queue configuration. The setting is ignored.
```

SASv9 GRID CONFIGURATION MIGRATION

Convert/Import LSF Configuration From Platform Web Services

```
$> ./sas-viya workload-orchestrator config grid import-from-swo \  
--url http://mySwoHost.mycompany.com \  
--user swoGridAdmin \  
--pwd swoGridAdminPwd \  
--validate-only
```