

Date: 3 January 2014

SAS® Visual Analytics Hardware 6.3 Sizing Guideline for Single Machine, 16 Core Environment with 256 GB of RAM

A hardware sizing guideline provides information on the performance characteristics of a certain hardware platform when running a specific piece of SAS software. It guides customers in evaluating whether or not the platform is appropriate for them. This hardware guideline provides an overview of the configuration for a single machine deployment of SAS Visual Analytics version 6.3. The guideline is to enable SAS customers to evaluate a single machine configuration's appropriateness for their environment. Numerous considerations ensuring a successful single-machine deployment that meets your business objectives follow.

Key Requirements:

1. SAS Visual Analytics 6.3 requires SAS 9.4. Other current SAS configurations don't need to be upgraded to SAS 9.4 but the SAS VA 6.3 piece requires this.
2. Windows is available for SAS VA 6.3 non-distributed (SMP) configurations only.
3. Non-distributed (SMP) SAS VA configurations can be virtualized.
4. Intel-Xeon E5 chips, 2670-2680-2690 for example, are required. The newer E5 2667v2 – 2680v2 -2690v2 and 2697v2 also work and increase memory up to 512 GB.

Sizing Considerations

Three key considerations for evaluating the appropriateness of the platform are:

1. Identifying the size of the largest individual data table to be loaded into RAM
2. Quantifying the total number of heavy users
3. Quantifying the total number of light users

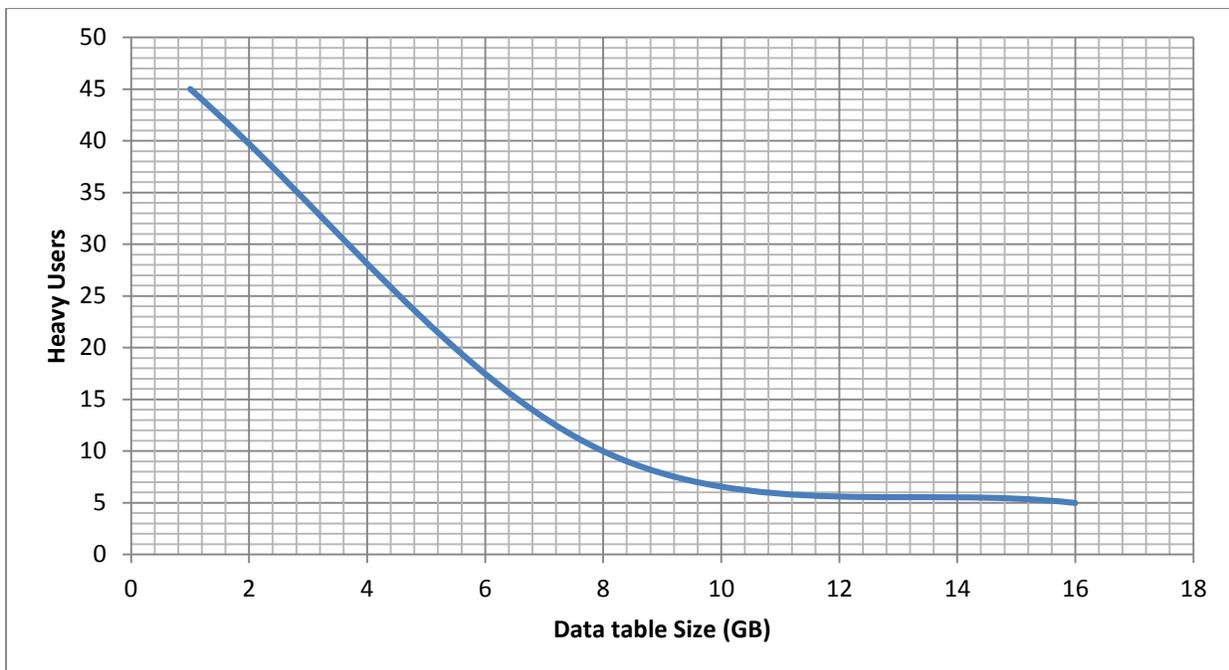
By our definition a heavy user is potentially any VAE users as well as others that run correlational analysis with multiple variables, Box plots with 4 or more measures and Cross tabs with 4 or more class variables each. The larger the data table the larger the impact.

A light user runs simple line or box chart with one or two class variables and up to five measures, bubble plots, heat maps or geo maps or a consumer of reports.

Performance Considerations

The following chart shows the relationship between the number of heavy users and the table size with the maximum number of heavy users on the x axis and the maximum in gigabytes on the y axis. Adding more users than the calculation allows may result in performance degradation for the end-user. To calculate the total number of light users this environment will support, take the following steps:

- Find the size of the largest table in the chart below. If the largest table size is larger than 16 GB, this environment is too small. See the Growth Considerations section for additional options.
- Read the chart to determine the maximum number of heavy users for that table size. If the number of heavy users exceeds the maximum identified in the table, this environment is too small. See the Growth Considerations section for additional options.
- Subtract the number of heavy users in the user base from the maximum number of heavy users in the chart.
- If the remaining number is greater than 0, multiply that number by 3 to identify the number of light users this environment can support. If the number of light users exceeds the number calculated, this environment is too small. See the Growth Considerations section for additional options.



For example, an 8 GB data table has a maximum of 10 heavy users. If you only have 5 heavy users, take the remaining 5 (10-5) and multiply by 3 (5 x 3= 15). The result is a total number of 5 heavy and 15 light users for this data table on the Single Machine instance of SAS VA.

Hardware Configuration

This Sizing Guideline is based on commodity hardware (Dell / HP or IBM) using the Intel-Xeon architecture. Other vendors that met the hardware specifications listed below can also be used.

The estimate uses the following baseline specifications:

SAS VA (Single Machine)	
# Nodes	1
CPU per node	2 x 8 core Intel Xeon E5-2670 series Chips (2.6 GHz)
RAM per node	256 GB
Memory Speed	1600 MHz
Operating System	Red Hat Enterprise Linux 6.4 or Windows x64
Local Disk per node	2 x 600 GB (10K RAID 1)
NIC	Dual 10 Gb NIC card

- Server power setting need to on the highest setting.
- Hyper-threading is recommended for all CPU's.
- The Text Analytics feature of VA 6.3 can impact overall system performance and aren't taken into consideration in this sizing.

Critical Data Considerations

1. This guideline does not address the data management resources necessary outside of SAS Visual Analytics. Getting data into SAS Visual Analytics and other ETL functions is solely the responsibility of the customer.
2. Maximum single largest table size to be loaded is 16 GB. Larger tables can be used but performance starts to degrade, i.e. slow down.
3. SAS VA can run with a single or multiple datasets loaded into memory
4. Up to 180 GB of tables can be loaded into memory
5. **Any changes to the hardware baseline or data specifications invalidates this sizing guideline**

Growth Considerations

What should a customer consider if individual table size limit is too small, but total tables stored is enough?

Performance of an individual table is directly related to the number of cores that the data is distributed across. If you need an individual table to be in the 20-30 GB range, but can limit your total tables size to ~200 GB, the customer may want to consider a single machine 32 core box.

However, if total tables size also needs to go up, you may want to consider moving to a multi-node environment. Please work with your Account Executive to have an official SAS Sizing executed.

What should a customer consider if total tables size needs to grow beyond the 90 GB limit?

Because of memory speed drop when moving from 256 GB to 512 GB on a single machine, we do not see a single machine as an option in this scenario. We would recommend the customer consider moving to a multi-node environment. Please work with your Account Executive to have an official SAS Sizing executed.

What should a customer consider if the number of users that a single machine can support isn't enough?

If reducing your single largest table size is an option, the guidelines above will help you figure out whether this environment can scale to meet your user needs.

A 32 core single machine may be an option if the number of users desired is slightly larger for the 16 core single machine guidelines. However, if the user count is significantly larger than the recommendations, a multi-node environment may give you more growth room for the future. Please work with your Account Executive to have an official SAS Sizing executed.

Disclaimer

Copyright © 2014, SAS Institute Inc., Cary, NC, USA, All Rights Reserved.

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration. Other brand and product names are trademarks of their respective companies.

This document is the confidential and proprietary property of SAS Institute Inc. It may contain approaches, techniques, and other information proprietary to SAS, and shall not be disclosed in whole or in part to third parties without the prior written consent of SAS.

The hardware sizing guideline herein is a no fee service. The guideline is based on expected usage and other technical details provided by and SAS' knowledge and expertise regarding their respective products. This hardware sizing recommendation is not a performance benchmark, and does not provide any performance guarantee. 's actual usage of the SAS software on the recommended hardware may not be consistent with the usage data provided by resulting in differing performance. SAS disclaims any liability with regards to (i) the hardware sizing recommendation provided herein; or (ii)'s reliance on the recommendation. acknowledges that it is responsible for all costs associated with procuring any recommended hardware and acknowledges that changes in hardware may result in SAS software upgrade license fees being payable.