

Server & IO Subsystem Considerations

Server Considerations:

Based on the single-tier topology of this environment it is recommended that at **least 8GB of RAM per CPU core is used**. This is based on a number of factors including topology, OS, amount of data, type of analyses, typically user behavior and learning over time. The justification for this request is as follows:

1. The SAS Metadata Server & Compute Server are residing on the same (virtual) machine and each utilizes resources intensely, but somewhat differently.
 - a. *Metadata Server*: This is equivalent to the Master Brain of the system. Everything that happens through SAS initially gets processed by this server. Definitions of things such as security, authentication, users/groups, servers, data access/definitions, scheduling, reports, etc. are all located within this server. Upon startup the metadata repository is loaded into RAM and stays there until it is shut down. The footprint of this repository is a function of how much metadata exists; hence it typically grows over time. We have seen repositories as small as 2GB all the way up to 50GB in size, depending on the types and scopes of implementations.
 - b. *Compute Server*: This server is the work-horse of the SAS implementation. Any interaction with the SAS environment performed by users occurs through this server. Things such as reading and manipulating data, creating graphs & reports, performing analyses, etc. are all run through this server. This intensity of resource utilization with this server is highly variable depending on what users are working on. In some cases with basic queries the utilization is no intense. In other cases where users are pulling a lot of data into SAS, manipulating it, creating graphical, tabular and statistical analyses on the data the resource utilization increases.
 - c. *User Behavior*: what is typically seen in every new environment is users utilize more resources than is typically needed. Users typically copy large amounts of data into SAS to perform analyses. As users get better acquainted with SAS they learn how to pass more of the query to the database to process and returning the results to SAS. This efficiency results in lowered server resource utilization. However, this utilization also begins to climb again when users begin to understand all of the things they are capable of with the environment.

Disk Performance Explanation and Considerations:

The main concern when setting up a file system is to ensure that SAS gets the sustained I/O bandwidth needed to complete the SAS jobs in the timeframe required by the SAS users.

- The I/O rate for SAS data varies by process.
- With the large, block I/O that SAS does, the I/O throughput rate is very important to SAS. The I/O throughput rate for SAS depends on what SAS tasks are being done: For most Extract, Transform, and Load (ETL) processes, 50-75MB/second of I/O throughput for each file system SAS is using is required per SAS session.
 - For end-user exploitation, 15-25MB/second per user per file system is required. However, for

heavy SAS users, 50-75MB/second per user per file system is required.

- Disks allocated to SAS working files - including "SASWORK" and "Web Report Studio Query Cache" - require 15-25MB/second per user for most SAS processes with heavy SAS users requiring 50-75MB/second per user.
- Disks allocated for SAS utility files for SAS threaded procedures and OLAP processing - "UTILLOC" - require 50-75MB/second per user.

Type	Sustained IO
ETL	50-75 MB/sec/Session
WORK (Temporary area) and Utility Files	50-75 MB/sec/Session
End users	15-75 MB/sec/Session

Make sure that you use the best file system for your operating system that follows the types of I/O your SAS