

The SAS System Client/Server Products and Solutions for the VSE Operating System

Paul Mortensen – SAS Institute Inc., Cary, North Carolina

Abstract:

Today's market offers customers many client/server products and options. The client/server product or model that will satisfy one application is not necessarily a good choice for your next set of client/server needs. What is needed is a product with strong client/server support and the flexibility to meet the varied needs of today's client/server applications and users.

The availability of SAS/CONNECT and SAS/SHARE for Release 6.08 of the SAS System for the VSE/ESA operating system provides VSE sites a flexible and powerful client/server tool set.

Why Client/Serve ?

With more Information System departments supporting multiple computing platforms with different machine architectures and data formats, and an increasing need by applications to access and process data regardless of location, the demand for client/server computing is growing fast.

The Gartner Group defines client/server as "one process requesting services or data from another process running on a separate machine". This definition is fairly generalized and as we will see later on, there are several methods or models for implementing a client/server solution for an application.

The five major reasons for implementing a client/server application are:

- data integration
- hardware independence
- balance processor workloads
- exploiting the strengths of different processors
- application integration.

Data integration is the ability to access data regardless of where it resides or what format it is in. A good client/server solution or product will allow an application to access and manage the data it needs across a variety of computing platforms and not limit the data formats that can be read.

Hardware independence allows you to develop and execute your application anywhere and to communicate between platforms. It requires that the client/server software provide a certain degree of portability between the platforms involved.

Processor workload can be balanced by distributing the work between the platforms, thereby relieving undue stress on any one processor and avoiding potential bottlenecks.

Many times a client/server solution is used to exploit the strengths of the various processors involved. Some processors have data intensive or compute intensive strengths, while others have presentation intensive strengths. This is an important consideration when planning the

distribution of the various components of an application between the client and the server.

Application integration gives you the ability to port applications across platforms and allows for the segmentation of applications. This requires a high degree of portability from the client/server software.

Why Client/Serve using VSE/ESA ?

There are several excellent reasons for using VSE/ESA as part of your client/server solution.

- VSE/ESA provides strong client/server support
- VSE/ESA has all the strengths of a mainframe system
- preservation of your VSE investment (hardware, software, and data)


Not too long ago, there was much speculation about the impending demise of the mainframe computers. The press was reporting that they would soon be replaced by the smaller faster desktop systems. "Downsizing" was the fashion and everyone was supposedly working out the best way to migrate their applications off the mainframes and on to these slick new micro wonders. The justification was that the desktop systems were cheaper to purchase and maintain, and offered as much, or more computing horse-power for a fraction of the cost of the mainframe systems.

Next, we started hearing about "right-sizing". Apparently there were some computing tasks that the mainframes were still better suited for. For applications that require data intensive processing, systems and data security, and data and resource centralization, many of the desktop systems still had a way to go to be considered a viable alternative.

The latest buzzword getting attention is "recentralizing". This is the shifting of those processing tasks that require centralized resources (data, peripherals, etc.) back on to a centralized processor. This latest trend, to me, is a reaffirmation of traditional strengths that the mainframe systems offer.

The VSE/ESA operating system provides strong client/server support with CICS, CICS/DDM, LANRES/VSE, ADISM, SQL/DS, Visalift and VisualGen, Open Connect System's TCP/IP, and the SAS System. VSE/ESA has all the strengths of a mainframe system, such as centralized data and processing, efficient transaction processing, systems management capabilities (reliability, security, data backup and restore), and support for unattended distributed operations.


According to several telephone surveys made by Intelliware Inc. in 1993, an overwhelming majority of mainframe sites, specifically VSE sites, also have some type of UNIX workstation and/or personal computers installed on site.



Many companies want to preserve their VSE investment in hardware, software, applications, and data. Full scale conversion is an expensive proposition.

Cost is another reason to consider using a mainframe system such as VSE/ESA in your

client/server application. In 1994, the Xephon Group did a study that showed that the five year average cost per user for server support on a mainframe was almost half of what server support costs were for PCs and UNIX.




Models for Client/Server Solutions

When considering the implementation of a client/server solution, it is very important to research and plan carefully. This section provides some information on several widely used models.

There are three basic components to any client/server solution, the client, the server, and the data. The client is the platform where the user or application is requesting data or services from. It is usually a PC or UNIX workstation, but another mainframe platform (MVS, VSE or VM/CMS) can act as a client as well. The server is the remote processor providing the requested data or services, and then there is the target data which can reside either on the client or the server.

Applications can generally be divided into four basic components. These are: the data, data access and management, application logic, and presentation. The data can reside in different places and be in various formats. The data access and management component are the engines i/o functions, or access methods that the application uses to read the data. The application logic component, is where the data is processed into useful information, and the presentation component is the final stage where the information generated by the application is presented to the user, usually in the form of graphs, reports, query screens, or EISs.

The Gartner Group has described a cooperative processing model that defines five separate types of client/server applications as shown in the figure below. The five model types are differentiated by the way that they distribute the application components between the client and the server. There are two data access models, one distributed function model, and two presentation models.



The two data access models differ from one another in that the distributed data access model splits the data access and management component between the client and the server, while the remote data access model keeps the data access and management all on the server side of the network dividing line. The application logic and presentation components are handled by the client for these models.

Similarly, the two presentation models differ in that the distributed presentation model splits the presentation component between the client and the server while the remote presentation model keeps the presentation all on the client side. With the presentation models, the data access and management and application logic components are placed on the server side of the network.

The distributed function model divides the application logic task between the client and the server.

There are many factors to consider when choosing which model to use for an application.

Resources available

- What hardware, software and data is currently available ?
- What resources do you anticipate having in the future ?

Workloads of each processor

- Which processors are currently stressed ?
- Which processors are under utilized ?

Network traffic

- How would each model affect network traffic in terms of data flow, and communications across the network ?

Security issues

- What security issues exist, especially in terms of data access ?

Data requirements

- Where is the data and what format is it in ?
- How current must the data be for each application ?

The SAS System and Client/Server

The SAS System provides a powerful and flexible toolset for developing client/server applications across a variety of platforms, including VSE/ESA, and addresses all the requirements that should exist in a vendor product for use in a client/server scenario. Those requirements being:

- Access to all data
- Connectivity of network
- Choice of client/server processing models
- Support of rapid application development
- Maintainability and extendibility

Access to all Data

The ability to access data regardless of its location or format is an important feature for any client/server toolset. The SAS System provides read and write data engines for transparent access to the following data sources and platforms, with new combinations added regularly.

VSE	SQL/DS VSAM VSE Librarian DL/I (4Q94)	OS/2	PC File Formats SQL Server AS/400 DB/2
MVS	ORACLE ADABAS DATACOM/DB DB2 IDMS/R IMS System 2000	WINDOWS	PC File Formats SQL Server ODBC
		VAX/OPEN VMS	ORACLE Rdb INGRES
CMS	ORACLE SQL/DS System 2000	UNIX	ORACLE SYBASE INGRES INFORMIX

Connectivity of Network

The ability to communicate across the network is critical to any client/server application. The SAS System supports all commonly used communications protocols for all the platforms supported (VSE, MVS, CMS, UNIX, OS/2, OpenVMS, Windows 3.x, and Windows NT). Terminal-based protocols such as 3270, Async, and HLLAPI, are supported as well as peer-to-peer protocols like APPC, TCP/IP, DECnet, and Netbios.

On each system, the SAS System can run stand-alone or cooperatively with the SAS System running on other processors, running the same operating system or a different operating system. For the VSE/ESA operating system, the SAS System supports the following communications protocols with other platforms:

OPERATING SYSTEM	COMMUNICATION PROTOCOL
IBM Mainframe: VSE, MVS, VM	APPC TCP/IP
UNIX	TCP/IP
OS/2	APPC, TCP/IP ASYNCR, HLLAPI
Windows	APPC, HLLAPI

Choice of Client/Server Processing Models

The SAS System provides a flexible and powerful toolset for building an Information Delivery System in any client/server solution. The SAS System's Multi-Vendor Architecture (MVA) allows not only the portability of applications across platforms, but also supports the segmentation of applications at any layer across platforms as well. This gives SAS users the ability to run each component of their client/server application on the platform where it makes the most sense.

Support of Rapid Application Development

The SAS System provides a number of features that assist in rapid application development

(RAD) across all platforms, including VSE/ESA. SAS reports, graphs, and programs can be easily generated using SAS/ASSIST software. With its graphical user interface, you can point and click your way through tasks and then save the generated SAS code for customization and future use. SAS/AF, SAS/FSP, and SAS/EIS allow rapid application development of GUI based applications.

Other factors that make the SAS System a good tool for rapid application development are:

portability of applications across platforms

the portability of SAS code allows for development on one platform and porting it to another

single vendor product interface

once you learn SAS on one platform, your skills and applications are readily transferrable to other platforms

selection of computing modes is interchangeable

traditional batch or interactive mode, using the SAS Multiuser interface for VSE

numerous host interfaces on VSE

VSE Librarian access,
DL/I access,
VSAM file access,
SQL/DS access,

reduces terminal switching on VSE

allowing for interactive submission of batch jobs,
VSE system console support,
VSE POWER queue management support,
dynamic file access/allocation for VSAM, EPIC, DYNAM-D, and sequential disk

Maintainability and Extensibility

One of the key factors for success when implementing a client/server solution is simplicity. You want to simplify your maintenance, problem determination and resolution, and future extendability as much as possible. When your client/server solution involves multiple vendors and products across multiple platforms, problem identification and resolution can become troublesome and complicated. With the SAS System, you greatly reduce the number of vendors and products involved across all of your client/serving platforms, which should, in turn, greatly simplify the vendor identification and problem resolution process.

The SAS System allows your client/server solutions to be adaptable for the future with little or no rewriting of applications. The portability of the SAS System insures that your application investment is secure no matter what platforms are added or removed from the network. As new platforms are introduced to the market, Multi-Vendor Architecture allows for a timely port of the SAS System to that platform.

The SAS System provides two sets of services that enable the development of client/server applications. The first is *compute services*, which give you the ability to access all of the computing services available on your network. The second is *remote data services*, which give you access to data regardless of where it is stored.

Compute Services

Compute services allow you to move any or all of the segments of your application from one processor to another. You can maximize the computing services within your organization this way by moving compute intensive operations to processors that have resources available to handle the job. The results from the remote execution are returned to the local SAS session, and

can then be used with further local processing, and/or presented by a local graphical user interface.

Remote compute services can also give you access to peripherals on other systems that may not be available to the local system, such as plotters, or printers.

Remote Data Services

Remote data services can be divided into two subgroups, data transfer services, and remote library services.

Data Transfer Services (DTS)

Data transfer services (DTS) provide a method for moving a copy of the data from one machine to another. Data can be moved in either direction, from the remote to the local processor, or the other way, whichever makes sense for the application's needs. SAS/CONNECT software provides support for DTS with the UPLOAD and DOWNLOAD procedures.

One of the major benefits of DTS is that it can be used to offload work from the remote system to one or more local processors, thereby improving response times for production systems executing in the remote environment. Once a copy, or subset copy of the data is moved to the local processor, applications can be run against it there without further impact or CPU consumption on the remote system.

Another benefit of DTS is that after the data has been copied locally, your application and its users will be insulated from any network or remote system performance degradations or failures.

There are some disadvantages to data transfer services that must be considered. Transferring data between systems causes multiple copies of that data to be created, and thereby incurs an increased use of DASD storage resources across the platforms.

There may be security issues related to the data that may restrict your ability to make multiple copies of that data on local (non secured) processors. In cases such as these, remote library services may be a better option.

For large volumes of data, the increased network traffic and the amount of time required to make the data transfer make be unacceptable to your application or your installation. You may need to consider transferring large volumes of data during off-peak hours as a work around.

With data transfer service, the data is only as current as the date and time of the last data transfer. If your application has a need for more current, or real-time data access, then perhaps remote library services is a better solution.

Remote library services (RLS)

Remote library services provides access to remote data as if it were in a local library. The data is not transferred between the systems, but instead remains in place on the remote system. This can allow multiple users to have concurrent access to the data while guaranteeing the integrity of the data. In the SAS System, the SAS/SHARE product provides the features that allow users to use remote library services.

RLS provides more transparent access to the data, the explicit steps of uploading or downloading the data are eliminated. There is only one copy of the data, reducing the storage resource overhead, allowing for more security control, and it gives users access to the data in

real-time.

The drawbacks to RLS are that it is not efficient for large amounts of data. Although a local copy of the data is not made, it must still be moved back and forth between the local and remote processors, which can have an adverse effect on the network traffic.

When multiple passes have to be made against the data, RLS performance can become an issue. It may be more advantageous to make a local copy of that data, hitting the network once, rather than paying the cost for each separate pass of the data.

When planning, careful thought must be given to the following issues: the amount of data to be accessed by the application, will single-user or multiple-user access to the data be necessary, will there be multiple passes made against the data, how will the network load affect your application, and vice-versa.

The SAS System and VSE/ESA

Many new VSE system specific functions have been added to release 6.08 of the SAS System for the VSE/ESA operating system.

SAS Multiuser system for interactive mode

- ICCF is no longer required,

- SAS Multiuser runs in a single partition and supports multiple users running on one image of the SAS System.

VSE Librarian access

- library/sublib members can be accessed and managed through the SAS System.

VSE System console support

- with proper authority, SAS Multiuser sessions can view the VSE system console and issue console commands.

Interactive submission of batch jobs

- VSE batch jobs can be submitted from within a SAS Multiuser session.

VSE/POWER queue management

- files in the RDR, LST, and PUN queues can be managed for your SAS Multiuser session.

Dynamic file access and allocation

- external JCL is no longer necessary to access files in sequential disk, VSAM/SAM managed space, or system managed space.

File security

- SAS Multiuser security table can be used to limit access to certain files and the use of VSE system interfaces on a user by user basis.

Conclusions

With the rising demand for connectivity between mainframe systems and other platforms, especially workstations, the need for client/server solutions are increasing as well. The key to successfully developing a client/server solution for your application lies in the ability to carefully segment your application, and implementing each component of the application on the platform where it will serve the needs of the application and the system best.

The SAS System offers a complete set of tools that provide the flexibility to implement the segments of your application where they best serve your needs. Compute services, data transfer services, and remote library services allow you to maximize resource usage, improve response times, and effectively increase enduser productivity.

SAS, SAS/AF, SAS/ASSIST, SAS/CONNECT, SAS/EIS, SAS/FSP, and SAS/SHARE are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries.

Other brand and product names are registered trademarks or trademarks of their respective companies.

References

Proceedings of the Nineteenth Annual SAS User's Group International Conference, The SAS System: A Complete Client/Server Solution, by Cheryl Garner, SAS Institute Inc., pages 247–257

Proceedings of the Nineteenth Annual SAS User's Group International Conference, Opening the Doors to VSE, by Tony Dean, SAS Institute Inc., pages 300–307

Techniques for Distributed Applications Development: The SAS System for VSE, a GUIDE presentation / June 1994, by Catherine Corbett, SAS Institute Inc.