

SAS® conversion considerations when migrating from HP OpenVMS VAX™ to HP OpenVMS Alpha™



Executive Summary	2
Introduction.....	2
SAS® versions on HP OpenVMS VAX™ and HP OpenVMS Alpha™	3
VAX D Floating Point to Alpha T Floating Point.....	3
Option 1.....	4
Option 2:.....	4
HOSTFMT= Option.....	5
SAS®/Toolkit Considerations.....	6
Other SAS® Data Files.....	7
SAS Data Views.....	7
SAS Catalogs.....	7
Conclusion.....	8
References	8
For more information.....	9

Executive Summary

This document communicates factors to consider and best practices for migrating SAS applications currently running on OpenVMS VAX to OpenVMS Alpha.

- SAS 8.2 is the last release of the SAS system for VAX. Existing VAX customers who desire to maintain their software upgrades must upgrade to SAS on the OpenVMS Alpha platform.
- Migrating from OpenVMS VAX to OpenVMS Alpha is the most straight forward conversion for the SAS System.
- A VMS cluster environment makes the process less difficult since the two can both exist simultaneously in a mixed architecture cluster. This allows SAS running on the Alpha system to read and convert the VAX data files directly without moving the files to a separate system. The OpenVMS cluster wide file system gives all nodes in the cluster equal access to files.

Introduction

The SAS System for OpenVMS Alpha first became available with SAS 6.09 in 1993. Alpha is based on Reduced Instruction Set Computer (RISC) architecture. The use of RISC architecture provided a significant performance boost for the SAS application software due to the nature of the application being CPU and memory-intensive based. The Alpha architecture is well suited for SAS due to its support for large amounts of memory, extremely fast processing speeds, and enhanced peripheral support solutions such as StorageWorks. All of these factors offer a significant performance improvement over VAX-based configurations.

Migration to an Alpha based processor from VAX systems may be desirable since VAX systems and peripheral components are no longer manufactured. While adequate stockpiles of VAX platform spares are available, the supply is limited and the accumulated support costs for stockpiled spares will continue to increase as stockpiles shrink. Ultimately, expenses for VAX systems will become far greater than support expenses for the newer, faster, and more reliable Alpha systems.

Some of the costs that need to be analyzed in making a decision of migrating to Alpha include:

- Hardware maintenance
- Facilities (i.e. electrical, A/C, footprint)
- Cost of new hardware
- Software license (cost difference)
- Cost of porting
- Need to take advantage of the latest release software enhancements

SAS® versions on HP OpenVMS VAX™ and HP OpenVMS Alpha™

This section outlines a brief history of SAS releases on OpenVMS VAX and OpenVMS Alpha. Also included is the minimum version of the OpenVMS operating system required for a particular release of the SAS system.

Table 1: SAS releases supported on OpenVMS VAX and OpenVMS Alpha

VAX Version	SAS Version	Release Date	Alpha Version	SAS Version	Release Date
5.0	6.08	1992 (many before this)	1.5	6.09	1993 (first release on Alpha)
5.4	6.09e	1995	6.1	6.12	Nov-96
6.2	7.0	Nov-98	6.2	7.0	Nov-98
6.2	8.0	Dec-99	7.1	8.0	Dec-99
6.2	8.1	Jun-00	7.1	8.1	Jun-00
6.2	8.2	Apr-01	7.1	8.2	Apr-01

VAX D Floating Point to Alpha T Floating Point

A major factor to consider when moving to the new RISC architecture is the storage of data types. All numeric variables within the SAS System are stored as double precision floating point. However, in Release 6.09 of the SAS System on OpenVMS Alpha, the data type for a numeric variable is IEEE double precision or T_FLOATING. The precision of a T_FLOATING value is expressed as 15 decimal digits, with a range of 1.7E+308 to 2.3E-308. This is different than the data type of numerical variables in the SAS System on OpenVMS VAX, where the data type of a numeric variable is D_FLOATING. The precision of a D_FLOATING value is 16 decimal digits, with a range of 1.7E+38 to 2.9E-29. There is not a D_FLOATING data type as part of the Alpha architecture; instead IEEE T_FLOATING is intended to take its place.

This means there is a potential for precision loss when moving from OpenVMS VAX to OpenVMS Alpha. For example, if a value is stored with a length of three bytes, the largest integer representation exactly on Alpha is 8,191 and on VAX it is 65,536. This means that moving a SAS data set which contains a variable named SCORE whose values include 8190 through 8199 and whose length is three bytes from VAX to Alpha would result in a data set whose values of SCORE are 8190, 8191, 8192, 8192, 8194, 8194, 8196, 8196, 8198, 8198 on Alpha. Notice that the values 8193, 8195, 8197 and 9199 could not be represented exactly in three bytes and truncating for storage in three bytes causes a loss in precision. Numeric precision issues are not unique to the OpenVMS operating system. One way to avoid the loss of precision is to increase the length of the variable SCORE in a DATA step on VAX before reading it into the SAS System for Alpha.

The default length of a numeric value stored in a SAS data set is eight bytes. Use a LENGTH statement to specify the number of bytes for storing numeric variables. This is useful (but not encouraged) when creating large permanent SAS data sets that require large amounts of disk space for storage. Numeric values that can be stored in less than eight bytes will help decrease the overall size of the data set and help conserve space. This conversion from the natural and default length of eight bytes will consume CPU cycles when doing the conversion. Under OpenVMS Alpha, the length of numeric values in the SAS System can range from three to eight bytes where the range under OpenVMS VAX is from two to eight bytes. When attempting to read a SAS data set of two bytes under OpenVMS Alpha, you will receive the following error:

```
ERROR: IEEE numbers with a length less than 3 are not supported. This data set contains
observations with numeric variables of length 2. The data set cannot be created/translated.
```

If production SAS data sets contain numeric values stored as two bytes (for any numeric stored in less than eight bytes this is recommended), there are two options for moving the data to the OpenVMS Alpha environment:

Option 1: The first option is to use the LENGTH statement in a DATA step on the VAX platform to change the length of two-byte numeric variables to three or more bytes before moving to the Alpha platform. Any attempt to change the length with a LENGTH statement in a DATA step while reading the data set on Alpha will not work. This is because variable length specified in the LENGTH statement affects the length of numeric variables in the output data set.

Option 2: The second option is to transport the SAS data set from VAX to Alpha. To protect against precision loss while moving between host systems, the XPORT (transport) engine adds one byte to any numeric value which has a length less than eight bytes. In other words, a numeric value stored in a length of two bytes will be read in as a three-byte numeric value after being transported to OpenVMS Alpha. Using PROC CPORT/CIMPORT to transport the data will also increase by one the length of numeric values with a length of less than eight bytes.

Numeric values in a SAS data set created on the OpenVMS VAX system subsequently read using OpenVMS Alpha will be converted, observation-by-observation, as the data set is read in. The base I/O engine on OpenVMS Alpha performs this automatically. Any overhead due to the internal conversion by the software emulation available on the OpenVMS Alpha system may cause performance degradation. The SAS Institute technical staff **does not** recommend routinely using SAS data sets interchangeably between OpenVMS platforms. By default, the SAS System creates a new data set on Alpha it uses T_FLOATING values for numeric. This requires that in a mixed Architecture VMS cluster the VAX SAS data sets should be kept separate from the Alpha SAS data sets. The SAS System on VAX cannot read Alpha created SAS data sets.

Missing values under the SAS System for OpenVMS VAX are represented by the 28 smallest D_FLOATING values. Although it is unlikely to occur in typical SAS programs, these values can still be calculated. With the SAS System under OpenVMS Alpha, IEEE Not-a-Number (NAN) values are used to represent missing values. An IEEE NAN, is an IEEE floating-point bit pattern that represents something other than a valid numeric value. These numbers cannot be calculated. There should be no difference in the sorting order of missing values or in their functionality.

Character values are stored using the ASCII collating sequence on either the VAX or Alpha OpenVMS platforms.

HOSTFMT= Option

A new option called HOSTFMT= was introduced in Release 6.08 of the SAS System for VAX. This option enables the creation of non-native data sets for read access on either hardware platform running under the OpenVMS operating system. A non-native SAS file is any file that is being accessed on a platform other than the one on which it was created. Using this option it is possible to create a SAS data set with the AXP host format on the OpenVMS VAX system and then read it on the OpenVMS Alpha system without any internal conversion. The HOSTFMT= option is valid for reading or writing, but not updating SAS data sets. The following code illustrates the use of the HOSTFMT= option to create a SAS data set with AXP host format on the OpenVMS VAX platform:

```
/* LIBNAME option */  
libname alphasib base '[mydir]' hostfmt=asp ;  
data alphasib.history;
```

or

```
/* data set option */  
libname alphasib base '[mydir]' ;  
data alphasib.history (hostfmt=asp) ;
```

The non-native data set can be read on the OpenVMS VAX system but cannot be updated there. If attempting to edit the data set ALPHALIB.HISTORY using PROC FSEDIT, for example, a note similar to the following will appear:

NOTE: Using FSBROWSE since ALPHALIB.HISTORY cannot be opened for update.

Under Release 6.09 and higher on Alpha, the data set ALPHALIB.HISTORY can be read or updated as shown in the following example:

```
libname alphasib base 'vaxnode::disk:[mydir]' ;  
proc print data=alphalib.history ;  
proc fsedit data=alphalib.history ;
```

Running PROC CONTENTS on ALPHALIB.HISTORY now shows that the V609 engine was assigned, the file format is 607, and the host format is AXP. The file format value is still 607 since the file format for data sets created by the BASE engine (V607, V608 or V609) has remained the same since Release 6.07.

Note: In Release 6.08, the host format value, which refers to a non-native data set, is ALPHA. In Release 6.09 and beyond, the host format value is AXP.

When copying data sets using HOSTFMT=VAX | AXP to be utilized by the respective platform, the following error will be received if an index or sort information is still associated with the data set being copied:

ERROR: Invalid or inconsistent mode flags for libref.dsname.DATA.

To circumvent this problem, remove the index and/or sort information using PROC DATASETS as shown in the following example:

```
/* Use to remove the index */
proc datasets library=<libref> ;
    modify <dsname> ;
    index delete <index name> ;
run ;

/* Use to remove sort information */
proc datasets library=<libref> ;
    modify <dsname> (sortby=_null_) ;
run ;
```

The index and/or sort information can be created again once the data sets are part of the new target native platform. In this instance, once the data sets are created on the OpenVMS Alpha environment the index and sort information can be created and used again.

SAS®/Toolkit Considerations

The Alpha architecture and VAX architecture are not binary compatible. Therefore, images compiled and linked on one system are not compatible with images created on the other. User written procedures, formats, informats, and functions created for OpenVMS VAX with the SAS/TOOLKIT product will need to be recompiled and relinked before being used with the SAS System for OpenVMS Alpha. Due to the differences between the two architectures, changes may also have to be made to the source code to make the code compatible with Alpha.

The SAS/TOOLKIT product and protocols must be used for functions intended to run on the newer OpenVMS Alpha platform, and all such code must be compiled and linked on the OpenVMS Alpha system. Refer to the SAS/TOOLKIT documentation for more information.

Other SAS® Data Files

SAS Data Views

A SAS data file with an extension of .SASEB\$VIEW may be a SAS DATA step view, a PROC SQL view, or a SAS/ACCESS view descriptor. These files do not actually contain data, but store information about the data for retrieval. This category of files is not transportable using PROC COPY or PROC CPORT/CIMPORT. It is recommended that all types of SAS views be recreated on the OpenVMS Alpha system. However, it is possible to use a PROC SQL view on Alpha that was created on VAX. Performance considerations may dictate whether PROC SQL views are recreated using data with the AXP file format or the SQL view is called to access data with the VAX file format.

A PROC SQL view can be created on OpenVMS VAX and successfully read on OpenVMS Alpha. The data that is accessed by the SQL view must be accessible from the Alpha node. The following code is an example of how to create a PROC SQL view:

```
libname alphasib base '[mydir]' ; /* no HOSTFMT= option */  
  
proc sql ;  
  
    Create view alphasib.bicycles as select * from alphasib.history ;  
  
quit ;
```

SAS DATA step views are not transportable, nor can HOSTFMT= option be used to create a non-native data step view on the OpenVMS VAX. A DATA step view is a collection of DATA step statements that are executed when the view is called. This does not support host-specific options. New DATA step views on the OpenVMS Alpha system must be created. Since source code must be saved when creating DATA step views, creating new DATA step views should involve copying the SAS source code to the OpenVMS Alpha and then rerunning the initial code to create new DATA step views.

SAS Catalogs

SAS catalogs created on the OpenVMS VAX platform must be converted before using them on the OpenVMS Alpha system. Use the CPORT procedure to create a transport file on the OpenVMS VAX system and the CIMPORT procedure to read the transport file on the Alpha platform. The transport file can be transferred from the VAX to Alpha platform via communications software such as DECnet or FTP (File Transfer Protocol), or by using the OpenVMS DCL COPY command. The COPY command would be the simplest way to create a transfer file when the VAX and Alpha systems are part of a mixed VMS cluster or where the VAX and Alpha systems share cluster wide storage.

The following examples are provided to illustrate the use of PROC CPORT and PROC CIMPORT in moving SAS catalogs from VAX to the Alpha platform:

On the VAX side, use PROC CPORT to create a transport data set that contains all of the catalogs in a particular directory called [MYPROJECT] will be copied into the transport file called PORTFILE.DAT. The following code illustrates one way to create the transport file:

```
/* Run on the VAX */  
  
libname toalpha 'vaxdisk1:[myproject]' ;  
  
proc cport library=toalpha file='[myproject]portfile.dat' memtype=catalog ;  
  
run ;
```

Running this SAS program code results in the creation of a transport data set called PORTFILE.DAT in the [MYPROJECT] directory. This transport file is ready to be transferred to the Alpha via communications software. The most important factor in successfully transferring a transport data set is to retain the internal format of the file. This generally means specifying a binary transfer.

The following DCL command issued from the Alpha system will copy the transport data set from VAX to Alpha:

```
$ COPY vaxnode::vaxdisk1:[myproject]portfile.dat [newproject]*
```

On the OpenVMS Alpha side, use PROC CIMPORT to read the transport data set into a native Alpha SAS catalog. The following code illustrates one way to do this:

```
/* Run of the Alpha */  
libname onalpha '[newproject]';  
proc cimport library=onalpha infile='[newproject]portfile.dat' memtype=catalog;  
run;
```

Error messages of "truncated record", "internal error from getting input data", and "given transport file is bad," indicate that the format of the transport file was modified during transfer and cannot be read properly. To fix this, transfer your transport file again verifying that no file attributes are changed so as to make sure it is a binary transfer. Other catalogs can be transported in a similar manner as described above. PROC REPORT catalog entries are not transportable

Conclusion

Moving applications from OpenVMS VAX to OpenVMS Alpha is straightforward and can be a simple process. In order to take advantage of the increase in performance and the newer features the RISC architecture has to offer, the better choice would be to move the production applications to OpenVMS Alpha to run under Release of 6.09 and higher of the SAS System. There are several factors to consider when moving to an OpenVMS Alpha architecture. These include storage of data types and what types of SAS files make up your applications. The SAS System provides several different techniques for converting your SAS applications to run under the OpenVMS Alpha system.

SAS data sets can be read and written to or from either OpenVMS platform. To ensure optimum performance, the HOSTFMT= option should be used when creating native SAS data sets on either the VAX or the Alpha system. For small or temporary jobs requiring data stored in a non-native data set, the SAS System will handle the conversion of numeric values automatically.

References

SAS Institute "Installation Instructions and System Manager's Guide for Release 8.2 of the SAS System under OpenVMS"

SAS Institute "SAS Companion for the OpenVMS Environment" Second Edition

HP OpenVMS "Performance Management"

HP OpenVMS "Upgrade and Installation Guide"

HP OpenVMS "System Manager's Manual: Essentials"

HP OpenVMS "System Manager's Manual: Tuning, Monitoring, and Complex Systems"

For more information

Document author - Carl Ralston an HP employee on site at SAS Institute in Cary, NC USA

email: Carl.Ralston@hp.com

phone: 919-531-5905

To learn more about upgrading from HP OpenVMS VAX to Alpha, please visit the following:

<http://h71000.www7.hp.com/advantage/index.html>

Finally, there are additional white papers, case studies, and configuration guides specific to HP and SAS on the HP/SAS partner website at:

<http://www.sas.com/partners/directory/hp/index.html>

© 2003 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

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

5982-3178EN, 11/2003