



**Dear Readers,**

The lazy days of summer have officially arrived – filled with vacations, holidays and more. While this may be the case for most of us as individuals, our companies keep on with business as usual. That's why I've filled this issue of the newsletter with articles and tips that will help you with your daily work.

You'll read about Live Web classes, an analytics lecture series, Microsoft Windows support, a book on mixed models and more. Plus, you'll get four – count 'em, four – technical tips and a white paper. Who says work slows down in the summer? Here at SAS, we're just getting started!

Here's to sun- and fun-filled days ahead...

A handwritten signature in black ink that reads 'Shelley Sessoms'. The signature is written in a cursive, flowing style.

**Shelley Sessoms**

Editor, *Your SAS Technology Report*

## Calculating Group Totals and the Counts Within Each Group

This example uses the SUM() function to sum the AMOUNT column, creating a new column named GRPTOTAL with a COMMA10. format.

The COUNT() function counts the number of occurrences of STATE within each group.

The GROUP BY clause collapses multiple rows for each group into one row per group, containing STATE, GRPTOTAL and the COUNT.

```
data one;
input state $ amount;
cards;
CA 7000
CA 6500
CA 5800
NC 4800
NC 3640
SC 3520
VA 4490
VA 8700
VA 2850
VA 1111
;
```

```
proc sql;
create table two as
select state
       ,sum(amount) as grptotal format=comma10.
       ,count(*) as count
from one
group by state;
```

```
quit;
```

```
proc print data=two;
run;
```

## Enhanced Printing Capability for SAS/IML<sup>®</sup> with the %PRINTIML Macro

SAS/IML is a powerful and flexible programming language capable of performing advanced mathematical and statistical computations. It is an integral component of the SAS intelligence architecture in that SAS/IML can read and write SAS data sets, just like the SAS DATA step.

One feature where SAS/IML could benefit in borrowing from the DATA step is its printing capability. The DATA step PUT statement is very advanced in options and has a multitude of features, while the SAS/IML PRINT statement is relatively simple in comparison. Ross Bettinger, a SAS Analytical Consultant, has written a macro (%PRINTIML, see Full Code tab) to remedy some of this deficiency. For example, the following code shows how %PRINTIML enhances SAS/IML's display of results:

```
%include 'printiml.sas' ;

proc iml ;

  reset fw=3 nocenter spaces=1 ;

  a = { 1 2 3, 4 5 6 } ;
  b = { 6 5 4, 3 2 1 } ;

  print '*** SAS/IML Printed Output ***' ;

  print a b ( a + b ) ;

  %PRINTIML( '%PRINTIML Printed Output' )

  %PRINTIML( a b ( a + b ) \ fw=3 spaces=1 )
quit ;
```

It produces the following output in the list file:

```
*** SAS/IML Printed Output ***

  A      B

  1  2  3  6  5  4  7  7  7
  4  5  6  3  2  1  7  7  7

%PRINTIML Printed Output

a =  1  2  3  b =  6  5  4  (a+b) =  7  7  7
    4  5  6      3  2  1      7  7  7
```

The %PRINTIML macro provides more labeling than the PRINT statement, especially when displaying complex expressions. Note that the result of evaluating the expression (a+b) is

displayed in the PRINT statement listing but is not labeled, while %PRINTIML clearly describes the result. Also, several printing-related RESET statement features that are not options for the PRINT statement are built into %PRINTIML.

Here is another example, involving row and column labeling:

```
%include 'printiml.sas';

proc iml ;

  reset fw=2 nocenter spaces=1 ;

  a = { 1 2 3, 4 5 6 } ;
  b = { 6 5 4, 3 2 1 } ;

  r_name = { "R1", "R2" } ;
  c_name = { "C1", "C2" } ;

  print '*** SAS/IML Printed Output ***' ;

  print a[ rowname=r_name colname=c_name format=1. ]
        b[ colname=c_name format=1. ]
        (a + b)[ colname=c_name format=1. ]
        ;

  %PRINTIML( '%PRINTIML Printed Output' )

  %PRINTIML( a[ rowname=r_name colname=c_name format=1. ]
            b[ colname=c_name format=1. ]
            ( a + b )[ colname=c_name format=1. ]
            \ fw=2 spaces=1
            )
quit ;
```

It produces the following output in the list file:

```
*** SAS/IML Printed Output ***

      B
A C1 C2 C3 C1 C2 C3 C1 C2 C3

R1 1 2 3 6 5 4 7 7 7
R2 4 5 6 3 2 1 7 7 7

%PRINTIML Printed Output

      C1 C2 C3      C1 C2 C3      C1 C2 C3
```

```
a = R1 1 2 3 b = 6 5 4 (a+b) = 7 7 7
    R2 4 5 6     3 2 1     7 7 7
```

Plainly, %PRINTIML produces results that are a little easier to understand than those produced with the PRINT statement.

### Using the %PRINTIML Macro

The %PRINTIML macro was designed to be used the same way that the PRINT statement is used, with the same inputs and additional RESET options. For comparison purposes, the PRINT statement syntax is:

```
PRINT < matrices >
  < (expression) >
  < "message" >
  < pointer controls >
  < [options] >
  ;
```

where the available options are:

```
colname=matrix
format=format
rowname=matrix
```

The %PRINTIML macro syntax is:

```
%PRINTIML( < matrices >
  < (expression) >
  < "message" >
  < pointer controls >
  < [options] >
  \ < RESET options >
)
```

where the available printing-related RESET options are described in Table 1:

**Table 1: %PRINTIML Reset Options**

Option Name	Default Value	Option Purpose
AUTONAME	Noautoname	Controls row and column labeling for matrices
CENTER	Nocenter	Specifies whether output is centered on page
FW	9	Sets the field width for printing numeric values

NAME	Noname	Specifies printing of the matrix name or label
SPACES	1	Specifies number of spaces between adjacent matrices

### A Final Note

The %PRINTIML macro uses the SAS/IML PRINT statement to display enhanced formatted results. %PRINTIML creates and embeds RESET statements before and after the PRINT statement. Because at the present time there is no way to restore the existing RESET settings before %PRINTIML is invoked, the default values in [Table 1](#) remain in effect after %PRINTIML terminates.

### Appendix

**Table A.1: Comparison Between SAS/IML RESET Statement and %PRINTIML Macro Options**

RESET Option	PRINT Statement	%PRINTIML Macro
AUTONAME	N	Y
CENTER	N	Y
FW	N	Y
NAME	N	Y
SPACES	N	Y

Note: Only options related to printing are considered.

**Table A.2: Comparison Between SAS/IML PRINT Statement and %PRINTIML Macro**

PRINT Statement Options	%PRINTIML Macro Options
Inputs	
Matrices	Matrices
(expression)	(expression)
'message'	'message'
Pointer controls	Pointer controls
Options	

COLNAME=matrix	COLNAME=matrix
FORMAT=format	FORMAT=format
ROWNAME=matrix	ROWNAME=matrix

---

**About the author:**

Ross Bettinger is a SAS Analytical Consultant. He provides support for Enterprise Miner and has been involved with data mining projects for 7 years. He has been a SAS user for 15 years. His professional interests are related to data mining, statistical analysis of data, feature selection and transformation, model building, and algorithm development.

## **ISO 9001:2000 Internal Auditor**

This Business Knowledge Series course is presented by Kim Bell.

This course is designed for those who want to obtain a thorough knowledge and understanding of the ISO 9001:2000 Standard. The course is suitable for participants from all backgrounds, including but not limited to manufacturing, healthcare, government, education, financial services, and general services. Individuals who will be participating, conducting, and/or managing an organization's internal audit process are especially encouraged to attend. This course is appropriate for individuals in both public and private sectors.

This course focuses on the following key areas: how the individual's roles and responsibilities contribute to the organization's quality management system, increasing revenue and market share through quality methods and timely customer response time, prioritizing improvement opportunities, measuring and analyzing data (including continual improvement data) within the organization, lowering costs and improving efficiency through optimum use of resources, determining and remaining focused on key improvement opportunities, optimizing costs and resources through a mutually beneficial supplier relationship, recognizing improvement opportunities within an organization, and auditing the individual's organization to promote quality and continual improvement efforts.

View the [complete outline](#) of this course. Live Web classes start at 1pm ET, 12pm CT, 11am MT, and 10am PT.

Read more at

[http://www.sas.com/apps/wtraining2/coursedetails.jsp?course\\_code=lwbiso&ctry=us](http://www.sas.com/apps/wtraining2/coursedetails.jsp?course_code=lwbiso&ctry=us)

## SAS® Support for Microsoft Windows Platforms

The table on this page provides an overview of the Microsoft Windows operating systems, both workstations and servers, that are supported under recent SAS releases, from SAS 8.2 to SAS 9.1.3. An "X" in the table indicates that the SAS release and Windows operating system combination is supported.

Included is a comment section that outlines which Microsoft Windows service packs are required or recommended and pertinent information related to [Windows x64](#). For more information about each release, see the system requirements for [SAS 8](#) and [SAS@9](#). SAS determines which Windows operating system its releases run under by following Microsoft's [Support Lifecycle Policy](#). When an operating system reaches its end of life, Microsoft reduces or ends support. SAS's [support lifecycle policy](#) defines the support available for SAS releases. To find out more about SAS 9.1.3, refer to the list of resources that are available on [SAS@9: The Highlights](#). You can find out more about the changes from SAS 8.2 to SAS@9 in [What's New?](#) More information is available about [SAS Foundation products](#) and [SAS Business Intelligence \(BI\)](#).

If you require more information that is not answered by this page or by system requirements, contact [SAS Technical Support](#).

Operating System	8.2	9.0	9.1	9.1.2	9.1.3	BI Capable	Comments
Windows 95	X						End of life per Microsoft's Support Lifecycle Policy
Windows 98	X						End of life per Microsoft's Support Lifecycle Policy
Windows Me	X						End of life per Microsoft's Support Lifecycle Policy
Windows NT Workstation 4.0	X	X	X	X	X	X	Recommend Microsoft Windows service pack 6a; end of life per Microsoft's Support Lifecycle Policy
Windows 2000 Professional	X	X	X	X	X	X	Recommend Microsoft Windows service pack 4 or higher
Windows XP Home Edition	X						
Windows XP Professional Edition	X	X	X	X	X	X	Recommend Microsoft Windows service pack 2 or higher
Windows XP Media Center Edition	X	X	X	X	X	X	Microsoft Windows service pack 2 or higher; follows same recommendations as XP Pro
Windows XP Tablet	X	X	X	X	X	X	Microsoft Windows service pack

PC Edition							2 or higher; follows same recommendations as XP Pro
Windows XP Professional x64					X		Only 32-bit Foundation SAS, shipped after September 2005
Windows XP Professional 64-bit Itanium Edition		X	X	X	X		Pulled from market by Microsoft and Hewlett-Packard
Windows NT Server 4.0, Standard Edition	X	X	X	X	X	X	Recommend Microsoft Windows service pack 6a; end of life per Microsoft's Support Lifecycle Policy
Windows NT Server 4.0, Terminal Server Edition	X	X	X	X	X	X	Recommend Microsoft Windows service pack 6a; end of life per Microsoft's Support Lifecycle Policy
Windows 2000 Server	X	X	X	X	X	X	Recommend Microsoft Windows service pack 4 or higher
Windows 2000 Advanced Server	X	X	X	X	X	X	Recommend Microsoft Windows service pack 4 or higher
Windows 2000 Datacenter Server	X	X	X	X	X	X	Recommend Microsoft Windows service pack 4 or higher
Windows Server 2003, Standard Edition (32-bit x86)	X	X	X	X	X	X	Recommend Microsoft Windows service pack 1 or higher
Windows Server 2003, Standard Edition R2 (32-bit x86)	X	X	X	X	X	X	
Windows Server 2003, Enterprise Edition (32-bit x86)	X	X	X	X	X	X	Recommend Microsoft Windows service pack 1 or higher
Windows Server 2003, Enterprise Edition R2 (32-bit x86)	X	X	X	X	X	X	
Windows Server 2003, Datacenter Edition (32-bit x86)	X	X	X	X	X	X	Recommend Microsoft Windows service pack 1 or higher
Windows Server 2003, Datacenter Edition R2 (32-bit x86)	X	X	X	X	X	X	

Windows Server 2003, Web Edition							Not supported: special-purposed Windows server operating system for use as a web server
Windows Small Business Server 2003	X	X	X	X	X	X	
Windows Storage Server 2003							Not supported: used for SAN/NAS support
Windows Server 2003, Enterprise Edition for Itanium-based System (64-bit)		X	X	X	X	X	Recommend Microsoft Windows service pack 1 or higher
Windows Server 2003, Datacenter Edition for Itanium-based System		X	X	X	X	X	Recommend Microsoft Windows service pack 1 or higher
Windows Server 2003 Standard x64 Edition					X		Only 32-bit Foundation SAS, shipped after September 2005
Windows Server 2003 Enterprise x64 Edition					X		Only 32-bit Foundation SAS, shipped after September 2005
Windows Server 2003 Datacenter x64 Edition					X		Only 32-bit Foundation SAS, shipped after September 2005
Windows Server 2003 R2 Standard x64 Edition					X		Only 32-bit Foundation SAS, shipped after September 2005
Windows Server 2003 R2 Enterprise x64 Edition					X		Only 32-bit Foundation SAS, shipped after September 2005
Windows Server 2003 R2 Datacenter x64 Edition					X		Only 32-bit Foundation SAS, shipped after September 2005

## **SAS Education Announces New Analytics Lecture Series**

This summer, SAS, the leader in analytics knowledge and technology, will launch a new 13-week analytics lecture series to discuss the latest trends and business issues in the industry. This new series will offer half-day educational sessions from some of the most highly regarded analytical thought leaders in the world. The program is delivered via SAS' state-of-the-art Live Web technology, so participants can interact with presenters and fellow attendees in real time.

**Read more** <http://support.sas.com/training/bks/als/>

## **SAS 2-for-1 Special (U.S. readers only)**

During the month of July, SAS Education invites you to take any of our U.S. public training courses and bring a co-worker for free! Space is limited. This offer is good only for students who register at the same time and attend the same course in July. Registration is now open!

**Read more** <http://support.sas.com/training/us/2for1.html>

## **Delivering the Fastest Time to Intelligence with Grid-Enabled SAS®**

Grid computing is quickly growing in importance as a way to harness the power of distributed computing resources. This white paper focuses on the evolution of grid computing with SAS and discusses the benefits it can bring to your organization, including increased performance, massive scalability and greater availability. The paper also includes several real-world examples of the power of SAS Grid Computing.

**Download the complete white paper at**

[http://www.sas.com/ctx/whitepapers/whitepapers\\_frame.jsp?code=343](http://www.sas.com/ctx/whitepapers/whitepapers_frame.jsp?code=343)

## ***SAS for Mixed Models, Second Edition***

**By:** Ramon Littell, George Milliken, Walter Stroup, Russell Wolfinger, and Oliver Schabenberger

**List price:** 89.95 USD

840 pages

**ISBN:** 1-59047-500-3

**Publisher:** SAS Press

**Publication Date:** February 2006

### **Description:**

The indispensable, up-to-date guide to mixed models using SAS®. Discover the latest capabilities available for a variety of applications featuring the MIXED, GLIMMIX, and NLMIXED procedures in this valuable edition of the comprehensive mixed models guide for data analysis, completely revised and updated for SAS®9. The theory underlying the models, the forms of the models for various applications, and a wealth of examples from different fields of study are integrated in the discussions of these models:

- random effect only and random coefficients models
- split-plot, multilocation, and repeated measures models
- hierarchical models with nested random effects
- analysis of covariance models
- spatial correlation models
- generalized linear mixed models
- nonlinear mixed models

Professionals and students with a background in two-way ANOVA and regression and a basic knowledge of linear models and matrix algebra will benefit from the topics covered. Includes a free CD-ROM with example SAS code!

**SAS Products Addressed:** Base SAS, SAS/ST

**Releases:** 9.1.3, 9.1.2, 9.1, 9.0, 8.2, 8.1, 8.01, 8.0

**Operating Systems:** 64-bit Enabled AIX, 64-bit Enabled HP-UX, 64-bit Enabled Solaris, ABI+ for Intel Architecture, AIX, HP-UX, HP-UX IPF, IRIX, LINUX X86\_64, Linux, Linux on Itanium, Microsoft Windows for IPF, OS/2, OpenVMS Alpha, OpenVMS VAX, Solaris, Tru64 UNIX, VM/CMS, WINDOWS X64, Windows, Windows NT Workstation, z

Order on-line today at

<http://www.sas.com/apps/pubscat/bookdetails.jsp?pc=59882&promo=EN>

## **We Want Your Ideas!**

Now is your chance to get an up-close look at all things SAS. The [Inside SAS Web seminar series](#) is your chance to experience SAS in totally new and different ways – ways that will help you enhance your job performance every day. You'll hear SAS experts discuss current products and technical topics in depth, and view a special segment that highlights hot SAS technology, including tips to help you drive innovation. Want to see a particular topic? Have an idea for a future Web seminar?

Contact [Natasha George](#) at [Natasha.George@sas.com](mailto:Natasha.George@sas.com) and let her know your thoughts and suggestions.

## **SAS' Ninth Annual Data Mining Conference Returns to Las Vegas**

### **October 23-24**

To help individuals learn the latest data mining theories, trends and best practices in the field, SAS hosts M2006, its ninth annual data mining conference, this October 24-25 in Las Vegas. The largest data mining conference of its kind, M2006 provides attendees a forum for exchanging ideas with hundreds of data mining practitioners and more than 40 of the most respected data mining [experts](#) in the world. We hope to see you there!

**Read more** <http://www.sas.com/events/dmconf/>

## DB2 Bulk Loading

If you have a lot of data to load into DB2 from SAS, the SAS/ACCESS Interface to DB2 offers the BULKLOAD=YES option. This will invoke the DB2 LOAD utility, enabling you to bulk load the rows of data as a single unit, which can significantly enhance performance for larger tables, where the speed of the bulk load process outweighs the overhead costs of the bulk load setup.

The SAS/ACCESS Interface to DB2 uses DSNUTILS, an IBM DB2 stored procedure that invokes the DB2 LOAD utility, to implement bulk loading. DSNUTILS is included in DB2 Version 6 and later. It is also available in DB2 Version 5 through a maintenance release. Because the LOAD utility is complex, verify with your database administrator that this utility is available and familiarize yourself with it before using it through SAS/ACCESS.

*Note:* When using the bulk loader, always look at the SYSPRINT output for information about the load.

### Preparing to load: the BLDB2TBLXST and BL\_DB2LDEXT options

The DB2 LOAD utility does not create tables; it loads data into existing tables. By default, SAS/ACCESS creates a table before loading data into it. If you want to invoke the utility for an existing table, specify BL\_DB2TBLXST=YES to tell the engine that the table already exists.

If instead the table does not exist, and you only want to generate control and data files for a subsequent create and load, specify BL\_DB2TBLXST=NO in conjunction with BL\_DB2LDEXT=GENONLY.

*Note:* Data to be loaded into an existing table must match the column types of the table.

SAS does not verify the input data against the table definition. Any incompatibilities are flagged by the LOAD utility.

### Configuring the load: the BL\_DB2LDCT1, BL\_DB2LDCT2, and BL\_DB2LDCT3 options

The BL\_DB2LDCT1 and BL\_DB2LDCT2 options allow you to pass parameters to the load process. By using these options, you can collect statistics about the table being loaded, compress its content, load data using a different character set, or avoid putting the table in copy-pending status when 'LOG NO' is also specified.

The parameters specified by using BL\_DB2LDCT1 are added to the load command after the LOAD keyword and before the 'INTO TABLE' clause. The parameters specified by using BL\_DB2LDCT2 are added to the load command between the table name and the column list. The parameters that you can pass depend on which release of DB2 you are running. See the IBM DB2 documentation for your release for the list of options you can pass by using BL\_DB2LDCT1 and BL\_DB2LDCT2.

BL\_DB2LDCT3 allows you to pass extra options that are added at the end of the control file. These options are not related to the load process, but they can be used to call other DB2 utilities, such as the REPAIR utility, to reset the tablespace status when 'LOG NO' is used and NOCOPYPEND is not specified. See the examples at the end of this section to see how this option can be used.

*Note:* Be sure to use uppercase when passing parameters to the LOAD utility with the BL\_DB2LDCT1, BL\_DB2LDCT2, and BL\_DB2LDCT3 options

### **Configuring the load: the BL\_DB2IN, BL\_DB2REC, and BL\_DB2PRINT options**

These options identify the file names for control statement input (BL\_DB2IN), table input (BL\_DB2REC) and report output (BL\_DB2PRINT). The default values for these options are respectively SYSIN, SYSREC, and SYSPRINT.

### **Input to the load (non-SMS): the BL\_DB2UNITCOUNT and BL\_DB2DEVT\_PERM options**

The SYSREC file contains the data to be loaded into the DB2 table. The default unit assignment for this file is SYSDA, but this can be overridden with the BL\_DB2DEVT\_PERM option. For large input files spanning multiple volumes, use the BL\_DB2UNITCOUNT option to indicate the number. The value specified for BL\_DB2UNITCOUNT must be an integer from 1 to 59. However, the maximum for a particular unit type may be less, so ask your storage administrator for this number.

*Note:* If the value specified with BL\_DB2UNITCOUNT exceeds the maximum number of volumes for the unit, an error is returned. The option is ignored if its value is greater than 59.

### **Input to/output from the load (SMS): the BL\_DB2DATACLAS, BL\_DB2MGMTCLAS, and BL\_DB2STORCLAS options**

The SAS/ACCESS Interface to DB2 allows SMS parameters to be specified for the SYSREC, as well as SYSIN and SYSPRINT, files that are allocated for the bulk load operation. The SMS data, management, and storage classes can be specified by using the BL\_DB2DATACLAS, BL\_DB2MGMTCLAS, and BL\_DB2STORCLAS options.

### **The BL\_DB2DATACLAS option**

Use this option to specify a data class for a new SMS-managed data set. The storage administrator at your installation defines the names of the data classes that you can specify with BL\_DB2DATACLAS. SMS ignores this option if you specify it for a data set that SMS does not manage. Also, if SMS is not installed or is not active, the operating system ignores any data class passed through BL\_DB2DATACLAS. This option applies to the control file (BL\_DB2IN), the input file (BL\_DB2REC), and the output file (BL\_DB2PRINT) for the bulk loader.

*Note:* For SMS managed data sets, the data class determines whether they can extend on multiple volumes. When BL\_DB2DATACLAS and BL\_DB2UNITCOUNT are both specified, the latter overrides the unit count values for the data class.

### **The BL\_DB2MGMTCLAS option**

Use this option to specify a management class for a new SMS-managed data set. The storage administrator at your installation defines the names of the management classes that you can specify with BL\_DB2MGMTCLAS. If SMS is not installed or is not active, the operating system ignores any management class passed through BL\_DB2MGMTCLAS. This option applies to the control file (BL\_DB2IN), the input file (BL\_DB2REC), and the output file (BL\_DB2PRINT) for the bulk loader.

### **The BL\_DB2STORCLAS option**

Use this parameter to specify a storage class for a new SMS-managed data set. The storage administrator at your installation defines the names of the storage classes that you can specify with BL\_DB2STORCLAS. The storage class contains the attributes that identify a storage service level to be used by SMS for storage of the data set. If SMS is not installed or is not active, the operating system ignores any management class passed through BL\_DB2STORCLAS. This option applies to the control file (BL\_DB2IN), the input file (BL\_DB2REC), and the output file (BL\_DB2PRINT) for the bulk loader.

*Note:* BL\_DB2STORCLAS replaces any storage attributes that are specified by using the BL\_DB2DEVT\_PERM option.

### **Restarting a failed load: the BL\_DB2RSTRT, BK\_DB2TBLXST and BL2LDEXT options**

In case of failure, you can restart a bulk load operation. Restart functionality is controlled by the BL\_DB2RSTRT option, which tells the LOAD utility whether the current load is a restart and, for a restart, indicates where to begin. When you specify a value other than NO, you must also specify BL\_DB2TBLXST=YES and BL\_DB2LDEXT=USERUN.

Valid values for this option are:

- NO: indicates a new invocation of the LOAD utility, not a restart. This is the default.
- CURRENT: indicates to restart at the last commit point.
- PHASE: indicates to restart at the beginning of the current phase.

How do these options interact with DB2? To answer this question, consider how the DB2 load utility works. The utility has 10 phases: UTILINIT, RELOAD, SORT, BUILD, SORTBLD, INDEXVAL, ENFORCE, DISCARD, REPORT, and UTILTERM. The ability to restart a load operation is determined as follows:

- If a failure occurs during the UTILINIT phase, restart the utility from the beginning.
- If a failure occurs during the RELOAD, ENFORCE or DISCARD phase, restart the utility using BL\_DB2RSTRT=CURRENT.

- If a failure occurs during the SORT, INDEXVAL, or REPORT phase, restart the utility using BL\_DB2RSTRT=PHASE.
- If a failure occurs during the BUILD phase, restart the utility using BL\_DB2RSTRT=PHASE if the REPLACE option was also used. If the load was appending data, the utility is not restartable. It must be terminated and the indexes must be rebuilt.
- If the utility fails during the UTILTERM phase, there is usually no need to restart it.

*Note:* When the load SORTKEYS option is used and the utility fails during the RELOAD, SORT, or BUILD phase, the value CURRENT or PHASE restarts the utility from the beginning of the RELOAD phase. For a complete description of all the phases of the DB2 load utility and their ability to be restarted, see the DB2 documentation from IBM.

## Examples

The following example creates a SAS data set and uses the SAS/ACCESS Interface to DB2 to load it into a DB2 table using the DB2 bulk load utility.

```

/* Create a small SAS sample data set */
data work.customers;
  input custname $ 1-10
         custnum
         custcity $ 16-30
datalines;
Beach Land 16 Ocean City
Coast Shop 3 Myrtle Beach
Coast Shop 5 Myrtle Beach
Coast Shop 12 Virginia Beach
Coast Shop 14 Charleston
Del Mar 3 Folly Beach
Del Mar 8 Charleston
Del Mar 11 Charleston
New Waves 3 Ocean City
New Waves 6 Virginia Beach
Sea Sports 8 Charleston
Sea Sports 20 Virginia Beach
Surf Mart 101 Charleston
Surf Mart 118 Surfside
Surf Mart 127 Ocean Isle
Surf Mart 133 Charleston

/* Open the SAS/ACCESS Engine to DB2 */
libname db2lib db2 ssid=db2a;

/* Load the SAS data set into a DB2 table */
data db2lib.customers (bulkload=yes
  dbtype=(custnum="smallint"));
  set work.customers;

```

```
run;
```

The following example uses the SAS/ACCESS Interface to DB2 to copy one DB2 table to another using the DB2 Bulk Load Utility:

```
libname db2lib db2 ssid=db2a;  
data db2lib.customers (bulkload=yes  
  bl_db2tblxst=yes  
  bl_db2ldct1='RESUME YEST');  
set db2lib.new-customers;  
run;
```

## Using FILENAME Statements to Access and Register External Files in SAS® ETL Studio

An external file is a file that is created and maintained by a host operating system or by another vendor's software application. SAS ETL Studio provides external file Source Designer wizards that enable you to register external files in a metadata repository. After an external file is registered, you can add it to a SAS ETL Studio job as a source or a target.

A fileref is a reference to FILENAME statement. In the external file Source Designer wizards, you can use a fileref instead of a filename and path to specify the physical location of the external file. You can give the fileref a simple and descriptive name that will be easier to work with than filenames and paths, which are often long and complicated.

The FILENAME statement can be specified in two main ways:

- add it as a pre-process to the job that will contain the metadata for the external file
- add it to the autoexec.sas file on the SAS application server that is being used to access the file.

**Note:** If you use a fileref instead of a physical path in an external file Source Designer, features that depend on interactive access to the file will not work unless you add the FILENAME statement to the appropriate autoexec.sas file. These features include verifying that the file exists and viewing the contents of the file.

If you cannot add a FILENAME statement to the autoexec.sas file on the SAS application server that is being used to access the file, then you can add it as a pre-process to the job that will contain the metadata for the external file. To add a FILENAME statement as a pre-process to a job, open the Properties window for the job and go to the **Pre and Post Processing** tab. Add the FILENAME statement to the **Pre-Processing** section.

Follow these steps to use a fileref in an external file Source Designer wizard:

1. From the SAS ETL Studio desktop, select **Tools** ➔ **Source Designer**.
2. From the Source Designer selection window, select the Delimited External File wizard or the Fixed Width External File wizard.
3. From the External File Location window, enter the fileref in the **File name** field. Click **Advanced** to access the Advanced File Location Settings window.
4. Click **File Name Quoting**, and select **No quotes around file name**. The fileref that is generated by your FILENAME statement cannot be read in the **File name** field of the External File Location window unless quotation marks are turned off.
5. Finish processing the external file that is referenced by the fileref.

## **Webcasts and Events**

### **[JMP® User Conference](#)**

**June 20-21**

**Cary, NC**

Attend exciting and insightful sessions, plus roundtable discussions, a scripting workshop, a Genomics Discovery event and exclusive training courses.

### **[Lafler's Upcoming Free Webinar](#)**

**June 27**

**12:30-1:30 p.m. ET**

SAS Press is hosting a free, live Webinar with Kirk Lafler on *Exploring Dictionary Tables with PROC SQL*.

### **[IT Management Summit on BI](#)**

**Multi-City Tour**

This series was created for busy IT professionals to acquire knowledge and insight from industry pundits as well as real-world experience on a variety of important topics concerning IT. Check the Web for dates and cities.

### **[Inside SAS: Spotlight on SAS® Enterprise Guide®](#)**

**On-Demand Web Seminar**

Now statisticians, business managers and programmers can become their own directors and producers! Find out how SAS Enterprise Guide integrates an array of analytics with the power of SAS in a friendly interface that allows business analysts to produce the reports they need.