



THE
POWER
TO KNOW.

SAS[®] Scalable Performance Data Server[®] 4.45 User's Guide

The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2008. *SAS® Scalable Performance Data Server® 4.45: User's Guide*. Cary, NC: SAS Institute Inc.

SAS® Scalable Performance Data Server® 4.45: User's Guide

Copyright © 2008, SAS Institute Inc., Cary, NC, USA

All rights reserved. Produced in the United States of America.

For a hard-copy book: No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, or otherwise, without the prior written permission of the publisher, SAS Institute Inc.

For a Web download or e-book: Your use of this publication shall be governed by the terms established by the vendor at the time you acquire this publication.

U.S. Government Restricted Rights Notice: Use, duplication, or disclosure of this software and related documentation by the U.S. government is subject to the Agreement with SAS Institute and the restrictions set forth in FAR 52.227-19, Commercial Computer Software-Restricted Rights (June 1987).

SAS Institute Inc., SAS Campus Drive, Cary, North Carolina 27513.

1st printing, May 2008

SAS® Publishing provides a complete selection of books and electronic products to help customers use SAS software to its fullest potential. For more information about our e-books, e-learning products, CDs, and hard-copy books, visit the SAS Publishing Web site at support.sas.com/publishing or call 1-800-727-3228.

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 registered trademarks or trademarks of their respective companies.

Table of Contents

Product Notes

- [SAS Scalable Performance Data \(SPD\) Server Product Notes](#)
-

SAS Scalable Performance Data (SPD) Server Usage

- [SAS Scalable Performance Data \(SPD\) Server Overview](#)
 - [Connecting to SAS Scalable Performance Data \(SPD\) Server](#)
 - [Accessing and Creating SAS Scalable Performance Data \(SPD\) Server Tables](#)
 - [Indexing, Sorting, and Manipulating SAS Scalable Performance Data \(SPD\) Server Tables](#)
 - [Using SAS Scalable Performance Data \(SPD\) Server with Other Clients](#)
 - [SAS Scalable Performance Data \(SPD\) Server Dynamic Cluster Tables](#)
-

SAS Scalable Performance Data (SPD) Server SQL Features

- [SAS Scalable Performance Data \(SPD\) Server SQL Planner](#)
 - [Connecting to the SAS Scalable Performance Data \(SPD\) Server SQL Engine](#)
 - [SAS Scalable Performance Data \(SPD\) Server SQL Planner Options](#)
 - [Parallel Join Facility](#)
 - [Parallel Group-By Facility](#)
 - [SAS Scalable Performance Data \(SPD\) Server STARJOIN Facility](#)
 - [SAS Scalable Performance Data \(SPD\) Server Index Scan](#)
 - [Optimizing Correlated Queries](#)
 - [Materialized Views](#)
 - [SAS Scalable Performance Data \(SPD\) Server SQL Extensions](#)
 - [Differences between SAS SQL and SAS Scalable Performance Data \(SPD\) Server SQL](#)
-

SAS Scalable Performance Data (SPD) Server SQL Reference

- [SAS Scalable Performance Data \(SPD\) Server SQL Syntax Reference Guide](#)
 - [SAS Scalable Performance Data \(SPD\) Server SQL Access Library API Reference](#)
-

SAS Scalable Performance Data (SPD) Server Reference

- [Optimizing SAS Scalable Performance Data \(SPD\) Server Performance](#)
 - [SAS Scalable Performance Data \(SPD\) Server Macro Variables](#)
 - [SAS Scalable Performance Data \(SPD\) Server LIBNAME Options](#)
 - [SAS Scalable Performance Data \(SPD\) Server Table Options](#)
 - [SAS Scalable Performance Data \(SPD\) Server Formats and Informats](#)
 - [SAS Scalable Performance Data \(SPD\) Server NLS Support](#)
-

SAS Scalable Performance Data (SPD) Server Frequently Asked Questions

- [SAS Scalable Performance Data \(SPD\) Server Frequently Asked Questions](#)

Copyright © 2008 by SAS Institute Inc., Cary, NC, USA. All rights reserved.

SAS Scalable Performance Data (SPD) Server Product Notes

- [Overview](#)
- [What's New in SAS Scalable Performance Data \(SPD\) Server 4.45?](#)
 - [SAS Scalable Performance Data \(SPD\) Server 4.45 Enhancements](#)
 - [CLUSTER LIST Statement Enhancements](#)
 - [Method Trace for SQL Enhancements](#)
- [What's New in SAS Scalable Performance Data \(SPD\) Server 4.44?](#)
- [What's New in SAS Scalable Performance Data \(SPD\) Server 4.43?](#)
 - [SAS Scalable Performance Data \(SPD\) Server 4.43 Enhancements](#)
 - [CLUSTER MODIFY Command for PROC SPDO](#)
 - [Table Option MEMNUM= for Dynamic Clusters](#)
- [What's New in SAS Scalable Performance Data \(SPD\) Server 4.42?](#)
 - [SAS Scalable Performance Data \(SPD\) Server 4.42 Enhancements](#)
 - [CLUSTER CREATE Option for Unique Indexes](#)
 - [Additional PROC SPDO Proxy Commands](#)
- [What's New in SAS Scalable Performance Data \(SPD\) Server 4.41?](#)
- [What's New in SAS Scalable Performance Data \(SPD\) Server 4.4?](#)
 - [SAS Scalable Performance Data \(SPD\) Server 4.4 Documentation](#)
 - [SAS Scalable Performance Data \(SPD\) Server 4.4 Platform Support Changes](#)
 - [New Platforms](#)
 - [Platforms No Longer Supported](#)
 - [SAS Scalable Performance Data \(SPD\) Server 4.4 and SAS Data Integration Studio](#)
 - [SAS Scalable Performance Data \(SPD\) Server 4.4 New Features](#)
 - [Materialized Views](#)
 - [SAS Scalable Performance Data \(SPD\) Server Profiling](#)
 - [LDAP Password Authentication](#)
 - [Dynamic Locking](#)
 - [Surfacing Ports through an Internet Firewall](#)
 - [SAS Scalable Performance Data \(SPD\) Server 4.4 Enhancements](#)
 - [MINMAX Table Indexing for Character Columns](#)
 - [Expression Support for STARJOIN](#)
 - [Dynamic Support for Larger Index Keys](#)

- [SORTEDBY Specification for Dynamic Clusters](#)
 - [Additional Backup, Restore, and List Options](#)
 - [Additional IXUTIL Options](#)

 - [What's New in SAS Scalable Performance Data \(SPD\) Server 4.3?](#)
 - [SAS 9.1.3 Compatibility and Large Table Support](#)
 - [SAS Scalable Performance Data \(SPD\) Server 4.3 and SAS 9.1.3 Password Encoding](#)
 - [SAS Scalable Performance Data \(SPD\) Server 4.3 and SAS Management Console](#)
 - [SAS Scalable Performance Data \(SPD\) Server 4.3 and SAS Data Integration Studio](#)
 - [SAS Scalable Performance Data \(SPD\) Server 4.3 Utility Requirements](#)
 - [SAS Scalable Performance Data \(SPD\) Server 4.3 SQL Planner Enhancements](#)
 - [SAS Scalable Performance Data \(SPD\) Server 4.3 MINMAX Table Indexing](#)
 - [SAS Scalable Performance Data \(SPD\) Server 4.3 WHERE Costing Improvements](#)
 - [SAS Scalable Performance Data \(SPD\) Server 4.3 Cluster Tables](#)
 - [SAS Scalable Performance Data \(SPD\) Server 4.3 Random Placement of Initial Data Partition Files in DATAPATH= Setting](#)
 - [SAS Scalable Performance Data \(SPD\) Server 4.3 Debugging Tools](#)
-

Overview

This document summarizes enhancements and changes in SAS Scalable Performance Data (SPD) Server 4.4, including the SAS Scalable Performance Data (SPD) Server 4.45, SAS Scalable Performance Data (SPD) Server 4.44, SAS Scalable Performance Data (SPD) Server 4.43, SAS Scalable Performance Data (SPD) Server 4.42 and SAS Scalable Performance Data (SPD) Server 4.41 maintenance releases. The enhancements and changes in SAS Scalable Performance Data (SPD) Server 4.3 are also included to provide users with a chronology.

The following compatibility information between the SAS System and the SAS Scalable Performance Data (SPD) Server 4.4 media is important:

- The SAS Scalable Performance Data (SPD) Server 4.4 CD-ROM includes client modules that are compatible with SAS 9.
 - SAS Scalable Performance Data (SPD) Server 4.4 is not compatible with SAS versions earlier than SAS 9. Refer to the appropriate SAS Scalable Performance Data (SPD) Server UNIX or Windows installation guide for more information about SAS software requirements for use with SAS Scalable Performance Data (SPD) Server 4.4.
 - For SAS 9.1.3 Service Pack 3 and earlier releases, you must rename the **sassqlu_for_sas913_sp3_and_earlier** modules from the SAS Scalable Performance Data (SPD) Server client installation to **sassqlu**. If you do not rename these modules for SAS 9.1.3 Service Pack 3 and earlier releases, problems will occur with SAS Scalable Performance Data (SPD) Server implicit pass-through SQL that uses three-part names. You will get an SQL parse error from SAS Scalable Performance Data (SPD) Server that causes the implicit pass-through SQL to fail.
-

What's New in SAS Scalable Performance Data (SPD) Server 4.45?

SAS Scalable Performance Data (SPD) Server 4.45, also called SAS Scalable Performance Data (SPD) Server 4.4 TSM5, is an interim release. SAS Scalable Performance Data (SPD) Server 4.45 contains maintenance fixes and feature enhancements that are not in SAS Scalable Performance Data (SPD) Server 4.44 and earlier releases.

The following feature enhancements are provided in SAS Scalable Performance Data (SPD) Server 4.45:

- SAS Scalable Performance Data (SPD) Server index performance after deleting values from a table has been improved. When values are deleted from a table, corresponding SAS Scalable Performance Data (SPD) Server 4.45 index values are virtually deleted, instead of physically deleted. Virtually deleted index values are invisible to SQL index queries. When you virtually delete an index value instead of physically deleting it, updating the index metadata requires less overhead resources. However, virtually deleted index values still occupy space in the disk image for the index. This disk image space can be reclaimed by using the ixutil index utility to reorganize the index. See the "SAS Scalable Performance Data Server Index Utility Ixutil" section of the online *SAS Scalable Performance Data Server 4.45: Administrator's Guide* for more detailed information about using the ixutil utility to generate virtually deleted value statistics and to reorganize indexes for optimum performance.
- SAS Scalable Performance Data (SPD) Server SQL has been enhanced to remove partial tables that can be created if certain SQL queries fail during execution. Previously, if an SQL statement such as, CREATE TABLE as SELECT failed during execution, partial tables were created that SAS Scalable Performance Data (SPD) Server could subsequently use, which led to erroneous results.

Note: The online *SAS Scalable Performance Data Server 4.45: Administrator's Guide* and the *SAS Scalable Performance Data Server 4.45: User's Guide* can be found at

<http://support.sas.com/documentation/onlinedoc/spds/index.html>.

SAS Scalable Performance Data (SPD) Server 4.45 Enhancements

SAS Scalable Performance Data (SPD) Server 4.45 introduces the following enhancements:

- [CLUSTER LIST Statement Enhancements](#)
- [Method Trace for SQL Enhancements](#)

CLUSTER LIST Statement Enhancements

The SAS Scalable Performance Data (SPD) Server CLUSTER LIST statement output lists the member tables of a dynamic cluster in numbered order. In SAS Scalable Performance Data (SPD) Server 4.45, more information is produced by the CLUSTER LIST statement.

Now, CLUSTER LIST statement output lists the member tables in a dynamic cluster, the maximum number of slots that are available in the cluster, and the unique index validation status.

The following example uses PROC SPDO to create a dynamic cluster that has a unique index, five

member tables, and space for a sixth table. Then, a CLUSTER LIST statement is issued. The output lists the dynamic cluster name, the names of each member table, the maximum number of member tables, and the unique index validation status.

```
PROC SPDO library=&libdom ;
CLUSTER CREATE natlsales ;
  mem=ne_region
  mem=se_region
  mem=central_region
  mem=nw_region
  mem=sw_region
MAXSLOT=6
UNIQUEINDEX=yes ;
```

NOTE: CLUSTER NATLSALES has been created with 6 maximum slots.

```
CLUSTER LIST natlsales;
```

```
Cluster Name NATLSALES, Mem=NE_REGION
Cluster Name NATLSALES, Mem=SE_REGION
Cluster Name NATLSALES, Mem=CENTRAL_REGION
Cluster Name NATLSALES, Mem=NW_REGION
Cluster Name NATLSALES, Mem=SW_REGION
```

NOTE: The maximum number of possible slots is 6.

NOTE: Unique index is validated in CLUSTER NATLSALES.

Method Trace for SQL Enhancements

The enhanced SAS Scalable Performance Data (SPD) Server method trace for SQL now includes the names of the source tables. Including source table names makes it easier to determine how the SAS Scalable Performance Data (SPD) Server SQL planner constructed the query in cases where multiple source tables were used.

To better understand the method trace enhancement, the following example shows a simple SQL statement that joins two tables and explicitly states the SAS Scalable Performance Data (SPD) Server execution methods that you want to use. The enhanced information content is displayed:

Before SAS Scalable Performance Data (SPD) Server 4.45:

```
PROC SQL;
CONNECT to sasspds(dbq= ... );
EXECUTE(RESET _method) by sasspds;
EXECUTE(CREATE TABLE c as SELECT t1.b t2.c
  FROM a t1, b t2
  WHERE t1.a = t2.a) by sasspds;
```

```
SPDS_NOTE: SQL execution methods chosen are:
sqxcrt
```


sqxjpll

SPDS_NOTE: Table X000003.C created, with 1 rows and 2 columns.

SAS Scalable Performance Data (SPD) Server 4.45:

```
PROC SQL;  
CONNECT to sasspds(dbq= ... );  
EXECUTE(RESET _method) by sasspds;  
EXECUTE(CREATE TABLE c as SELECT t1.b t2.c  
FROM a t1, b t2  
WHERE t1.a = t2.a) by sasspds;
```

SPDS_NOTE: SQL execution methods chosen are:

```
sqxcrta  
sqxjpll  
sqxrc ( [X0000001].A (alias = t1) )  
sqxrc ( [X0000001].B (alias = t2) )
```

SPDS_NOTE: Table X0000001.C created, with 1 rows and 2 columns.

See the section "Important SAS Scalable Performance Data (SPD) Server SQL Planner Options" in the online *SAS Scalable Performance Data Server 4.45: User's Guide* for more information on the SAS Scalable Performance Data (SPD) Server SQL planner.

[What's New in SAS Scalable Performance Data \(SPD\) Server 4.44?](#)

SAS Scalable Performance Data (SPD) Server 4.44, also called SAS Scalable Performance Data (SPD) Server 4.4 TSM4, is an interim release. SAS Scalable Performance Data (SPD) Server 4.44 contains maintenance fixes and feature enhancements that are not in SAS Scalable Performance Data (SPD) Server 4.43 and earlier releases.

The following feature enhancement is provided in SAS Scalable Performance Data (SPD) Server 4.44:

- **LASTCLUSTERMEMBER** is a new argument that you can use with the **MEMNUM=** table option to when you want to query or read from the last member table of a dynamic cluster. Instead of counting cluster members to determine the number (*n*) of the last member to use in the statement **MEMNUM=*n***; you can specify **MEMNUM=LASTCLUSTERMEMBER**. When you specify **MEMNUM=LASTCLUSTERMEMBER**, SAS Scalable Performance Data (SPD) Server selects the last member for you. No numeric value for *n* is required when you use the **LASTCLUSTERMEMBER** argument.

For example, to view the contents of the last member table in a cluster called **Clustername**, issue the statement:

```
PROC CONTENTS data=Clustername(MEMNUM=LASTCLUSTERMEMBER) ;
```

See the section "Querying and Reading Member Tables in a Dynamic Cluster" in the online *SAS Scalable Performance Data Server: User's Guide* for more information.

Note: The online *SAS Scalable Performance Data Server: Administrator's Guide* and the *SAS Scalable Performance Data Server: User's Guide* can be found at

<http://support.sas.com/documentation/onlinedoc/spds/index.html>

What's New in SAS Scalable Performance Data (SPD) Server 4.43?

SAS Scalable Performance Data (SPD) Server 4.43, or SAS Scalable Performance Data (SPD) Server 4.4 TSM3, is an interim release. SAS Scalable Performance Data (SPD) Server 4.43 contains maintenance fixes and feature enhancements that are not found in SAS Scalable Performance Data (SPD) Server 4.42 and previous releases.

The following feature improvements are provided in the SAS Scalable Performance Data (SPD) Server 4.43 release:

- The Windows 64 client WIA64 is supported. For more information on Windows 64 client support, see the "Installing and Configuring SAS Scalable Performance Data (SPD) Server Clients" topic in the "SAS Scalable Performance Data (SPD) Server Windows Installation Guide" documentation, or the "Configuring SAS Scalable Performance Data (SPD) Server Client Software" topic in the "SAS Scalable Performance Data (SPD) Server UNIX Installation Guide" documentation in the Installation section of the online *SAS Scalable Performance Data Server: Administrator's Guide*.
- The SAS Scalable Performance Data (SPD) Server Index Utility adds an option to provide index distribution statistics. For more information on the index distribution statistics option, see the "Ixutil Options" topic in the "SAS Scalable Performance Data (SPD) Server Index Utility Ixutil" documentation in the System Management section of the online *SAS Scalable Performance Data Server: Administrator's Guide*.
- Documentation has been added for using SAS Scalable Performance Data (SPD) Server with an Internet firewall. For more information on configuring SAS Scalable Performance Data (SPD) Server for use with Internet firewalls, see the "Using SAS Scalable Performance Data (SPD) Server with an Internet Firewall" documentation in the Security section of the online *SAS Scalable Performance Data Server: Administrator's Guide*.
- Internal performance optimizations have been implemented in SAS Scalable Performance Data (SPD) Server pass-through SQL. The updated SAS Scalable Performance Data (SPD) Server pass-through SQL improves the performance of queries that select a small number of columns from a table. The performance gains are most noticeable when joining tables where the result set contains only a small proportion of the total number of columns that exist in the joined tables.
- BY-clause control has been enhanced. A new server parameter option, [NO]BYINDEX, and a corresponding server macro, SPDSNBIX=, are used to permit BY-clauses on an indexed variable to sort the table using the variable's index. For more information on the [NO]BYINDEX server parameter option, see the "Setting Up SAS Scalable Performance Data (SPD) Server Parameter Files" documentation in the Configuration section of the online *SAS Scalable Performance Data Server: Administrator's Guide*. For more information on the SPDSNBIX= server macro, see the "SAS Scalable Performance Data (SPD) Server Macro Variables" documentation of the "SAS Scalable Performance Data (SPD) Server Reference" section of the online *SAS Scalable Performance Data Server: User's Guide*.
- SAS Scalable Performance Data (SPD) Server 4.43 features secure LDAP authentication for Solaris, AIX, HP-UX, and HP Itanium. For more information on secure LDAP authentication, see the "SAS Scalable Performance Data (SPD) Server Parameter File Configurations for LDAP" topic in the "Setting Up SAS Scalable Performance

Data (SPD) Server Parameter files" documentation in the Configuration section of the online *SAS Scalable Performance Data Server: Administrator's Guide*.

- The SAS Scalable Performance Data (SPD) Server cluster add operation has been modified to allow cluster reads while a cluster add is progress. The reader sees the state of the cluster when the read begins. Cluster adds still are exclusive operations; a cluster can only have one add operation occurring at any time.

Note: The online *SAS Scalable Performance Data Server: Administrator's Guide* and the *SAS Scalable Performance Data Server: User's Guide* can be found at

<http://support.sas.com/documentation/onlinedoc/spds/index.html>.

SAS Scalable Performance Data (SPD) Server 4.43 Enhancements

- [CLUSTER MODIFY Command for PROC SPDO](#)
- [Table Option MEMNUM= for Dynamic Clusters](#)

CLUSTER MODIFY Command for PROC SPDO

The PROC SPDO command set for dynamic clusters features a new CLUSTER MODIFY cluster command. The CLUSTER MODIFY command sets a MINMAXVARLIST attribute on one or more variables that belong to an existing dynamic cluster. When the SAS Scalable Performance Data (SPD) Server runs the CLUSTER MODIFY command, the dynamic cluster is unclustered while the variable modifications are made to the individual member tables. The cluster is recreated after the MINMAXVARLIST changes are completed.

For more information on the CLUSTER MODIFY command, see the "Modify Dynamic Cluster Tables" topic in the "SAS Scalable Performance Data (SPD) Server Dynamic Tables" document in the "SAS Scalable Performance Data (SPD) Server Usage" section of the online *SAS Scalable Performance Data Server: User's Guide*.

Table Option MEMNUM= for Dynamic Clusters

Dynamic clusters provide a table option MEMNUM=. The MEMNUM= option allows you to perform query or read operations on a single member table that belongs to the cluster. When you use the MEMNUM= option, SAS Scalable Performance Data (SPD) Server opens only the specified member table instead of opening all of the member tables that belong to the cluster.

For more information on the MEMNUM= option, see the "Querying and Reading Member Tables in a Dynamic Cluster" topic in the "SAS Scalable Performance Data (SPD) Server Dynamic Tables" document

in the "SAS Scalable Performance Data (SPD) Server Usage" section of the online *SAS Scalable Performance Data Server: User's Guide*.

What's New in SAS Scalable Performance Data (SPD) Server 4.42?

SAS Scalable Performance Data (SPD) Server 4.42 (or SAS Scalable Performance Data (SPD) Server 4.4 TSM2) is an interim release. SAS Scalable Performance Data (SPD) Server 4.42 contains maintenance fixes and feature enhancements that are not found in SAS Scalable Performance Data (SPD) Server 4.41 and earlier releases.

The following feature enhancements are provided in the SAS Scalable Performance Data (SPD) Server 4.42 release:

- The SQL RESET option, PRINTLOG, logs SQL queries to the SAS Scalable Performance Data (SPD) Server log. For more information, see the section "Important SAS Scalable Performance Data (SPD) Server SQL Planner Options," in the chapter, "Scalable Performance Data (SPD) Server SQL Features," in the *SAS Scalable Performance Data Server: User's Guide*.
- SQL LIBNAMEs and record-level locking LIBNAMEs are supported. For more information, see the section "LIBNAME Proxy Commands," in the chapter "SAS Scalable Performance Data (SPD) Server Operator Interface Procedure (PROC SPDO)," in the *SAS Scalable Performance Data Server: Administrator's Guide*.
- An SPD Management Proxy Manager utility is part of the SAS Management Console. The SPD Management Proxy Manager utility monitors SAS Scalable Performance Data (SPD) Server LIBNAME activity. For more information, see the section "Proxy Manager," in the chapter, "Administering and Configuring SAS Scalable Performance Data (SPD) Server Using the SAS Management Console," in the *SAS Scalable Performance Data Server: Administrator's Guide*.
- The SAS Scalable Performance Data (SPD) Server STARJOIN facility offers an IN-SET join strategy. The IN-SET join strategy allows you to use star schema processing when the star schema's fact table and dimension tables have simple indexes on join columns. For more information, see the section "SAS Scalable Performance Data (SPD) Server STARJOIN Optimization," in the chapter, "SAS Scalable Performance Data (SPD) Server STARJOIN Facility," in the *SAS Scalable Performance Data Server: User's Guide*.
- A BY clause sort optimization is available for cluster tables if the member tables in the star schema are sorted by the BY clause. For more information, see the section "Dynamic Cluster BY Clause Optimization," in the chapter, "SAS Scalable Performance Data (SPD) Server Dynamic Cluster Tables," in the *SAS Scalable Performance Data Server: User's Guide*.
- Secure LDAP authentication is available for Solaris and AIX. For more information, see the section "SAS Scalable Performance Data (SPD) Server Parameter File Configurations for LDAP," in the chapter, "Setting Up SAS Scalable Performance Data Server Parameter files," in the *SAS Scalable Performance Data Server: Administrator's Guide*.

The *SAS Scalable Performance Data Server: User's Guide* and *SAS Scalable Performance Data Server: Administrator's Guide* can be viewed from the following URL:

<http://support.sas.com/documentation/onlinedoc/spds/index.html>

SAS Scalable Performance Data (SPD) Server 4.42 Enhancements

- [CLUSTER CREATE Option for Unique Indexes](#)
- [Additional PROC SPDO Proxy Commands](#)

CLUSTER CREATE Option for Unique Indexes

The CLUSTER CREATE command in PROC SPDO has a new option that allows the user to specify whether unique indexes that are defined in the member tables should be validated and marked as unique in the cluster. If the UNIQUEINDEX option is set to No, then unique indexes are not validated, and the cluster metadata does not mark the indexes as unique within the cluster. If the UNIQUEINDEX option is not specified, then it defaults to YES and the indexes are validated and marked unique within the cluster.

The usage syntax for the CLUSTER CREATE command is:

```
CLUSTER CREATE clustername  
MEM=member_table1  
MEM=member_table2  
...  
MEM=member_table_n  
MAXSLOT=n  
UNIQUEINDEX=<yes/no>;
```

For more information on PROC SPDO commands, see the "SAS Scalable Performance Data (SPD) Server Operator Interface Procedure (PROC SPDO)" documentation in the System Management section of the online *SAS Scalable Performance Data Server: Administrator's Guide*.

Additional PROC SPDO Proxy Commands

The existing PROC SPDO command set has new commands that capture proxy information about pass-through SQL librefs. The new commands, LIST USERS/LOCKING and SET USER/LOCKING, capture information about record-level locking proxies that are associated with pass-through SQL librefs.

The new privileged OPER command OPER INTERRUPT enables certain users to interrupt long-running jobs. The new privileged OPER command OPER DISCONNECT drops the

proxy from its client. The OPER HALT and OPER RESUME commands are no longer supported.

For more detailed information on PROC SPDO proxy commands, see the "SAS Scalable Performance Data (SPD) Server Operator Interface Procedure (PROC SPDO)" documentation in the System Management section of the online *SAS Scalable Performance Data Server: Administrator's Guide*.

What's New in SAS Scalable Performance Data (SPD) Server 4.41?

SAS Scalable Performance Data (SPD) Server 4.41 (or SAS Scalable Performance Data (SPD) Server 4.4 TSM1) is an interim release. SAS Scalable Performance Data (SPD) Server 4.41 contains maintenance fixes and feature enhancements that are not found in SAS Scalable Performance Data (SPD) Server 4.4 and earlier releases.

The following feature enhancements are provided in the SAS Scalable Performance Data (SPD) Server 4.41 release:

- Indexes can be created on materialized views. For more information, see the section, "Materialized Views" in the chapter, "SAS Scalable Performance Data Server SQL Features," in the *SAS Scalable Performance Data Server: User's Guide*.
- The std, avg, stderr, uss, css, and var GROUP BY functions are supported for use with fast index scans. All functions that can use index scans can use the DISTINCT function as well. For more information, see the section, "SAS Scalable Performance Data (SPD) Server Index Scan," in the chapter "SAS Scalable Performance (SPD) Data Server SQL Features," in the *SAS Scalable Performance Data Server: User's Guide*.
- The SPDSBKUP utility backs up MINMAXVARLIST information and table column metadata such as FORMAT and LABEL. The SPDSRSTR utility restores the MINMAXVARLIST information and the table column metadata. For more information about SPDSBKUP and SPDSRSTR, see the chapter "SAS Scalable Performance Data Server Backup and Restore Utilities," in the *SAS Scalable Performance Data Server: Administrator's Guide*.
- When you create a sorted table using the ORDER BY clause with the CREATE TABLE SQL statement, the ORDER BY column in the new table is marked as sorted. Subsequent queries on the table that include an ORDER BY clause on the column will not cause the table to be re-sorted.

What's New in SAS Scalable Performance Data (SPD) Server 4.4?

- [SAS Scalable Performance Data \(SPD\) Server 4.4 User's Guide and Administrator's Guide](#)
- [SAS Scalable Performance Data \(SPD\) Server 4.4 and SAS Data Integration Studio](#)

- [SAS Scalable Performance Data \(SPD\) Server 4.4 Platform Support Changes](#)
- [SAS Scalable Performance Data \(SPD\) Server 4.4 New Features](#)
- [SAS Scalable Performance Data \(SPD\) Server 4.4 Enhancements](#)

SAS Scalable Performance Data (SPD) Server 4.4 Documentation

The *SAS Scalable Performance Data Server: User's Guide* and *SAS Scalable Performance Data Server: Administrator's Guide* have been removed from the SAS Scalable Performance Data (SPD) Server installation media. The documentation is available online at

<http://support.sas.com/documentation/onlinedoc/spds/index.html>.

Having the SAS Scalable Performance Data (SPD) Server 4.4 documentation on support.sas.com enables you to access it via your Web browser and facilitates rapid distribution of SAS documentation updates between successive SAS Scalable Performance Data (SPD) Server releases.

SAS Scalable Performance Data (SPD) Server 4.4 Platform Support Changes

New Platforms

SAS Scalable Performance Data (SPD) Server 4.4 has added support for the UNIX Solaris x64 platform.

Platforms No Longer Supported

SAS Scalable Performance Data (SPD) Server 4.4 no longer supports the Linux IA-64 platform or the HP Tru64 UNIX platform.

SAS Scalable Performance Data (SPD) Server 4.4 and SAS Data Integration Studio

You can integrate the processing power of SAS Scalable Performance Data (SPD) Server 4.4 with SAS Data Integration Studio. The plug-in file that SAS Scalable Performance Data (SPD) Server uses to integrate with the SAS Management Console can also incorporate SAS Scalable Performance Data (SPD) Server resources into the SAS Data Integration Studio user interface.

To incorporate SAS Scalable Performance Data (SPD) Server 4.4 functionality into the SAS Data Integration Studio user interface, copy the SAS Scalable Performance Data (SPD) Server 4.4 plug-in file

into the SAS Data Integration Studio **plugins** subdirectory.

The SAS Scalable Performance Data (SPD) Server 4.4 plug-in file is located at:

```
SASROOT/spdssmc/sas.smc.SpdsMgr.jar
```

Note: **SASROOT** represents the path to the base directory of the SAS software installation on your client machine. **spds44** represents the installed SAS Scalable Performance Data (SPD) Server software directory. The name of the installed SAS Scalable Performance Data (SPD) Server software directory varies according to the specific version and release of your SAS Scalable Performance Data (SPD) Server software. For example, the path to your SAS Scalable Performance Data (SPD) Server plug-in file might begin with **SASROOT/spds44**, **SASROOT/spds44tsm1**, or **SASROOT/spds44tsm2**, depending on whether you have the original SAS Scalable Performance Data (SPD) Server 4.4 software, or the first or second maintenance release of the SAS Scalable Performance Data (SPD) Server 4.4 software.

Copy the SAS Scalable Performance Data (SPD) Server 4.4 plug-in file to the SAS Data Integration Studio **plugins** directory:

```
SASROOT/SASETLStudio/9.1/plugins/sas.smc.SpdsMgr.jar
```

SAS Scalable Performance Data (SPD) Server 4.4 New Features

- [Materialized Views](#)
- [SAS Scalable Performance Data \(SPD\) Server Profiling](#)
- [LDAP Password Authentication](#)
- [Dynamic Locking](#)
- [Surfacing Ports through an Internet Firewall](#)

Materialized Views

A materialized view saves the results of a VIEW statement in a temporary SAS Scalable Performance Data (SPD) Server table. When the view is queried the temporary table is used for the query, instead of the entire view. If any of the input tables that comprise the view are modified, the materialized view dynamically updates the temporary table. A materialized view is supported only through the SAS Scalable Performance Data (SPD) Server SQL pass-through facility. A materialized view can result in significant performance improvements for queries that query the view.

For more information on materialized views, see the section, "Materialized Views," in the chapter, "Optimizing SAS Scalable Performance Data Server Performance," in the *SAS Scalable Performance Data Server: User's Guide*.

SAS Scalable Performance Data (SPD) Server Profiling

The SAS Scalable Performance Data (SPD) Server Process Profiler utility monitors and logs the activity of the SAS Scalable Performance Data (SPD) Server processes. Once the information is logged, the output can be formatted to be read into a SAS table for analysis.

The SPD Management Server utility in the SAS Management Console connects to the SAS Scalable Performance Data (SPD) Server Performance Server to provide real-time feedback of SAS Scalable Performance Data (SPD) Server process activity. The SAS Scalable Performance Data (SPD) Server 4.4 process profile panel dynamically refreshes SAS Scalable Performance Data (SPD) Server process activity such as memory and CPU usage. An SAS Scalable Performance Data (SPD) Server process is identified by its process ID (PID), and, if it is a proxy process, the SAS Scalable Performance Data (SPD) Server user name that is associated with the proxy is included.

This feature is available only for SAS Scalable Performance Data (SPD) Server 4.4 (and later) installed on UNIX.

For more information on SAS Scalable Performance Data (SPD) Server profiling, see the section, "SPD Process Profiler" in the chapter, "Administering and Configuring SAS Scalable Performance Data (SPD) Server Using the SAS Management Console," in the *SAS Scalable Performance Data Server: Administrator's Guide*.

LDAP Password Authentication

LDAP authentication causes SAS Scalable Performance Data (SPD) Server to authenticate a user password via LDAP, rather than the password in the PSMGR database. LDAP authentication allows an SAS Scalable Performance Data (SPD) Server user to have the same user name and password as the UNIX/Windows user name and password, if the UNIX/Windows logon user name and password meets the SAS Scalable Performance Data (SPD) Server user name and password character restrictions.

The administrator can select the mode of password authentication with server parameters; either via the PSMGR database or LDAP. Once selected all authentication will be done in that mode. With LDAP Authentication, a SAS Scalable Performance Data (SPD) Server user must still be entered in the SAS Scalable Performance Data (SPD) Server PSMGR database to maintain other information necessary for SAS Scalable Performance Data (SPD) Server, such as the user's groups and access level.

This feature is available only for SAS Scalable Performance Data (SPD) Server 4.4 installed on Solaris, AIX, HP-UX, HP Itanium, and Windows.

For more information on SAS Scalable Performance Data (SPD) Server LDAP authentication, see the section, "The Password Manager Utility," in the chapter, "Managing SAS Scalable Performance Data (SPD) Server Passwords, Users, and Table ACLs," in the *SAS Scalable Performance Data Server: Administrator's Guide*.

[Dynamic Locking](#)

Dynamic locking provides more flexible locking semantics on a domain, which allows multiple clients to share both Read and Write access to tables in the domain without getting locking failures. Dynamic locking differs from SPD record-level locking in that clients using dynamic locking connect to a separate SAS Scalable Performance Data (SPD) Server user proxy process for each LIBNAME connection in the domain. With record-level locking, all users share the same record-level locking proxy process. Having separate SAS Scalable Performance Data (SPD) Server user proxy processes lessens the chance of resource limits, and removes a single record-level locking point of failure for the record-level proxy process.

Dynamic locking can provide better performance than record-level locking in cases where concurrent reads and updates to a table are required, but the performance benefit needs to be measured on a case-by-case basis.

For more information on SAS Scalable Performance Data (SPD) Server dynamic locking, see the chapter, "Accessing and Creating SAS Scalable Performance Data Server Tables," in the *SAS Scalable Performance Data Server: User's Guide*.

[Surfacing Ports through an Internet Firewall](#)

SAS Scalable Performance Data (SPD) Server uses a client/server relationship, which means that the client cannot exist on the same host as the server. If the network environment has an Internet firewall, you have to control the ports that the SAS Scalable Performance Data (SPD) Server server and client use for communication so that those ports can be surfaced through the Internet firewall. Certain ports that the SAS Scalable Performance Data (SPD) Server uses are defined at start-up, and can therefore be easily controlled. However, ports are dynamically allocated to support each connection to the SAS Scalable Performance Data (SPD) Server and the subsequent user proxy processes that are created as a result of the connection. These ports are usually allocated as any available port. The MINPORTNO and MAXPORTNO server parameters are fully supported features in SAS Scalable Performance Data (SPD) Server 4.4. You can use the MINPORTNO and MAXPORTNO server parameters to control the dynamic ports that SAS Scalable Performance Data (SPD) Server uses.

For more information on surfacing ports through an Internet firewall, see the chapter, "Setting Up SAS Scalable Performance Data (SPD) Server Parameter Files," in the *SAS Scalable Performance Data Server: Administrator's Guide*, and the questions, "How do SAS Scalable Performance Data (SPD) Server client and server processes communicate?" and "How do I know which ports must be surfaced through an Internet firewall?" in the chapter, "SAS Scalable Performance Data (SPD) Server Frequently Asked Questions," in the *SAS Scalable Performance Data Server: User's Guide*.

[SAS Scalable Performance Data \(SPD\) Server 4.4 Enhancements](#)

- [MINMAX Table Indexing for Character Columns](#)
 - [Expression Support for STARJOIN](#)
 - [Dynamic Support for Larger Index Keys](#)
 - [SORTEDBY Specification for Dynamic Clusters](#)
 - [Additional Backup, Restore, and List Options](#)
 - [Additional IXUTIL Options](#)
-

MINMAX Table Indexing for Character Columns

The SAS Scalable Performance Data (SPD) Server table option for MINMAXVARLIST= has been enhanced to support character columns. The SAS Scalable Performance Data (SPD) Server WHERE clause planner uses the MINMAXVARLIST parameter values for a table to quickly determine whether a WHERE clause on the character column can be quickly evaluated as trivially true or false.

For more information on MINMAX table indexing, see the section, "MINMAX Indexes," in the chapter, "Optimizing SAS Scalable Performance Data Server Performance," in the *SAS Scalable Performance Data Server: User's Guide*.

Expression Support for STARJOIN

The SAS Scalable Performance Data (SPD) Server STARJOIN optimization has been enhanced to accept some queries that previously could not use the optimization. Those queries met the STARJOIN requirements, but the selected column was an expression, instead of a simple column. As a result, STARJOIN was not allowed.

For more information on expression support for STARJOIN, see the chapter "SAS Scalable Performance Data (SPD) Server STARJOIN Facility," in the *SAS Scalable Performance Data Server: User's Guide*.

Dynamic Support for Larger Index Keys

The SAS Scalable Performance Data (SPD) Server Indexes dynamically support an index key up to 32,608 bytes. An index key is the sum of the length of all of the columns that comprise the index. Previously, you had to reconfigure the BTREE_PAGESIZE server option to support larger index keys. With dynamic sizing of the index metadata to support larger index keys, reconfiguring the BTREE_PAGESIZE server option is no longer necessary and the option is now obsolete.

[SORTEDBY Specification for Dynamic Clusters](#)

SAS Scalable Performance Data (SPD) Server supports the SORTEDBY specification for columns that are defined in a dynamic cluster. To use the SORTEDBY specification, each member table in the dynamic cluster must have SORTEDBY specification set for the column. You set the SORTEDBY specification on a dynamic cluster in the same way you set it for a simple table.

```
PROC DATASETS library=libdomain;  
modify clustername(sortedby=<var>);  
quit;
```

The SORTEDBY specification assumes that the dynamic cluster was created using member tables that were added in the correct SORTEDBY order.

[Additional Backup, Restore, and List Options](#)

The SAS Scalable Performance Data (SPD) Server backup utility has added a -V option to provide verbose output. The -v option will log the full name of the backup file and table of contents file.

The SAS Scalable Performance Data (SPD) Server backup utility has added a -PROJ <dir> option to support backing up files in a domain project directory.

The SAS Scalable Performance Data (SPD) Server restore utility has added a -PROJ <dir> option to support restoring files to a domain project directory.

The SAS Scalable Performance Data (SPD) Server list utility has added an -S option to include the size (in bytes) of the component files that are listed.

The SAS Scalable Performance Data (SPD) Server list utility has added an -INFO option to get table information for a domain, including the number of component metadata, data, and index files for a table, and the accumulated size of the component files for a table.

For more information, see the chapter, "SAS Scalable Performance Data (SPD) Server Backup and Restore Utilities," in the *SAS Scalable Performance Data Server: Administrator's Guide*.

[Additional IXUTIL Options](#)

The SAS Scalable Performance Data (SPD) Server IXUTIL utility has added the -CREJIDX option to create a join index, the -DELJIDX option to delete a join index, the -STATJIDX option to print join index statistics, and the -LSTJIDX option to list the join indexes in a domain.

For more information, see the chapter, "SAS Scalable Performance Data (SPD) Server

What's New in SAS Scalable Performance Data (SPD) Server 4.3?

The enhancements and changes for SAS Scalable Performance Data (SPD) Server 4.3 are included to provide users with a chronology for the SAS Scalable Performance Data (SPD) Server feature set over the most recent releases.

- [SAS 9.1.3 Compatibility and Large Table Support](#)
- [SAS Scalable Performance Data \(SPD\) Server 4.3 and SAS 9.1.3 Password Encoding](#)
- [SAS Scalable Performance Data \(SPD\) Server 4.3 and SAS Management Console](#)
- [SAS Scalable Performance Data \(SPD\) Server 4.3 and SAS Data Integration Studio](#)
- [SAS Scalable Performance Data \(SPD\) Server 4.3 Utility Requirements](#)
- [SAS Scalable Performance Data \(SPD\) Server 4.3 SQL Planner Enhancements](#)
- [SAS Scalable Performance Data \(SPD\) Server 4.3 MINMAX Table Indexing](#)
- [SAS Scalable Performance Data \(SPD\) Server 4.3 WHERE Costing Improvements](#)
- [SAS Scalable Performance Data \(SPD\) Server 4.3 Cluster Tables](#)
- [SAS Scalable Performance Data \(SPD\) Server 4.3 Random Placement of Initial Data Partition Files in DATAPATH= Setting](#)
- [SAS Scalable Performance Data \(SPD\) Server 4.3 Debugging Tools](#)

SAS 9.1.3 Compatibility and Large Table Support

SAS Scalable Performance Data (SPD) Server 4.3 is compatible with the improved I/O infrastructure of SAS 9.1.3.

SAS Scalable Performance Data (SPD) Server 4.3 provides on-disk structures that are compatible with SAS 9 and the large table capacities that it supports. Enterprise-wide data mining often creates immense tables. In order to generate business intelligence quickly, the ability to update tables that contain billions of rows is more important than ever. Earlier versions of SAS Scalable Performance Data (SPD) Server were based on 32-bit architecture that supported just over 2 billion rows and 32,768 columns. SAS Scalable Performance Data (SPD) Server 4.3 is based on a 64-bit architecture that supports tables with over 9 quintillion rows and over 2 billion columns.

The architectural differences between SAS 9 and earlier SAS versions mean that SAS Scalable Performance Data (SPD) Server 4.3 cannot access SAS Scalable Performance Data (SPD) Server 3.x stores, and vice versa. For more information on sharing SAS Scalable Performance Data (SPD) Server 3.x and SAS Scalable Performance Data (SPD) Server 4.3 data stores, see the chapter, "SAS Scalable Performance Data (SPD) Server 3.x and SAS Scalable Performance Data (SPD) Server 4.4 Compatibility" in the *SAS Scalable Performance Data Server: Administrator's Guide*.

SAS Scalable Performance Data (SPD) Server 4.3 and SAS 9.1.3 Password Encoding

SAS Scalable Performance Data (SPD) Server 4.3 supports the integration of the SAS 9.1.3 PROC PWENCODE. This procedure permits scripts to be generated that do not explicitly contain secure passwords that could easily be used without authorization. You must run PROC PWENCODE in Base SAS software to enable the usage of script password encoding within SAS Scalable Performance Data (SPD) Server 4.3. See the Base SAS software documentation for detailed instructions on running PROC PWENCODE for use with SAS Scalable Performance Data (SPD) Server 4.3.

The following example shows an SAS Scalable Performance Data (SPD) Server 4.3 LIBNAME statement that uses the password encoding option:

```
libname mylib sasspds 'spdsdata'  
      server=kaboom.5200  
      user='spdsuser'  
      password='{sas001}c3BkczEyMw==';
```

SAS Scalable Performance Data (SPD) Server 4.3 and SAS Management Console

SAS Management Console is a Java application that provides a single point of control for managing multiple SAS application resources. Rather than using a separate administrative interface for each application in your enterprise intelligence environment, you can use the SAS Management Console interface to perform the administrative tasks that are required to create and maintain an integrated environment.

SAS Management Console manages resources and controls by creating and maintaining metadata definitions for entities such as:

- server definitions
- library definitions
- user definitions
- resource access controls
- metadata repositories
- job schedules

After installing the SAS Scalable Performance Data (SPD) Server 4.3 Java plug-in file, SAS Scalable Performance Data (SPD) Server administrators can use the SAS Scalable Performance Data (SPD) Server Server Manager utility in SAS Management Console to configure SAS Scalable Performance Data (SPD) Server 4.3 user and group passwords and ACLs, instead of using the traditional SAS Scalable Performance Data (SPD) Server

PSMGR database and PROC SPDO commands.

By default, SAS Management Console looks for plug-ins in the **plugins** subdirectory of each installed SAS product. The plug-in file that makes the SAS Scalable Performance Data (SPD) Server Server Manager utility available in SAS Management Console is located at:

SASROOT/spds43/plugins/sas.smc.SpdsMgr.jar

Note: **SASROOT** represents the path to the base directory of the SAS software installation on your client machine. The previous plug-in file path in the example is specifically for SAS Scalable Performance Data (SPD) Server 4.3. The plug-in file for SAS Scalable Performance Data (SPD) Server 4.4 is in [a different location](#).

SAS Scalable Performance Data (SPD) Server 4.3 and SAS Data Integration Studio

You can integrate the processing power of SAS Scalable Performance Data (SPD) Server 4.3 with other SAS software, such as SAS Data Integration Studio. The same Java plug-in file that SAS Scalable Performance Data (SPD) Server uses to integrate with SAS Management Console can be used to integrate SAS Scalable Performance Data (SPD) Server resources with the SAS Data Integration Studio user interface.

SAS Data Integration Studio enables data warehouse specialists to create and manage metadata objects that define sources, targets, and the sequence of steps for the extraction, transformation, and loading of data into data marts or warehouses. SAS Scalable Performance Data (SPD) Server can be an excellent tool for managing the large tables of data associated with large data marts and warehouses.

By default, SAS Data Integration Studio looks for plug-ins in the **plugins** subdirectory of the SAS Data Integration Studios installation. To incorporate SAS Scalable Performance Data (SPD) Server 4.3 functionality with the SAS Data Integration Studio user interface, copy the SAS Scalable Performance Data (SPD) Server 4.3 Java plug-in file into the SAS Data Integration Studio **plugins** subdirectory.

The SAS Scalable Performance Data (SPD) Server 4.3 Java plug-in file is located at:

SASROOT/spds43/plugins/sas.smc.SpdsMgr.jar

Note: **SASROOT** represents the path to the base directory of the SAS software installation on your client machine. **spds43** represents the installed SAS Scalable Performance Data (SPD) Server software directory. The name of the installed SAS Scalable Performance Data (SPD) Server software directory varies according to the specific version and release of your SAS Scalable Performance Data (SPD) Server software. For example, the path to your SAS Scalable Performance Data (SPD) Server Java plug-in file might begin with **SASROOT/spds43**, **SASROOT/spds43tsm1**, or **SASROOT/spds43tsm2**, depending on whether you have the original SAS Scalable Performance Data (SPD) Server 4.3 software, or the first or

second maintenance release of the SAS Scalable Performance Data (SPD) Server 4.3 software.

Copy the SAS Scalable Performance Data (SPD) Server 4.3 Java plug-in file to the SAS Data Integration Studio **plugins** directory:

```
SASROOT/SASETLStudio/9.1/plugins/sas.smc.SpdsMgr.jar
```

SAS Scalable Performance Data (SPD) Server 4.3 Utility Requirements

SAS Scalable Performance Data (SPD) Server 4.3 provides NLS (National Language Support) functionality for multiple languages and character sets in database operations. As a result, all SAS Scalable Performance Data (SPD) Server 4.3 utilities require access to the `<installdir>/bin64` directory, and you must ensure that the `<installdir>/bin64` directory is specified in your SAS Scalable Performance Data (SPD) Server 4.3 path statement.

Here is an example of a statement that specifies the necessary path:

```
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:InstallDir/bin64

export LD_LIBRARY_PATH
```

SAS Scalable Performance Data (SPD) Server 4.3 SQL Planner Enhancements

SAS Scalable Performance Data (SPD) Server 4.3 includes SQL planner optimizations. SQL planner optimizations improve the performance of frequent query types that are used in data mining solutions such as SAS Marketing Automation. A key enhancement to the SAS Scalable Performance Data (SPD) Server 4.3 SQL planner is optimizing correlated queries through query rewrite techniques. Correlated queries are common in business and analytic intelligence data mining. Another key enhancement is the tighter integration of the parallel GROUP BY technology in the SQL planner. Tighter integration adds performance benefits to nested GROUP BY syntax.

- [SAS Scalable Performance Data \(SPD\) Server 4.3 STARJOIN Facility](#)
- [SAS Scalable Performance Data \(SPD\) Server 4.3 Index Scans](#)
- [SAS Scalable Performance Data \(SPD\) Server 4.3 Optimized Correlated Queries](#)
- [SAS Scalable Performance Data \(SPD\) Server 4.3 Parallel GROUP BY](#)
- [SAS Scalable Performance Data \(SPD\) Server 4.3 Parallel Join](#)

SAS Scalable Performance Data (SPD) Server 4.3 STARJOIN Facility

The SAS Scalable Performance Data (SPD) Server 4.3 enhanced SQL planner includes the new STARJOIN facility. The SAS Scalable Performance Data (SPD) Server 4.3 STARJOIN facility validates, optimizes, and executes SQL queries on data that is configured in a star schema. Star schemas are composed of two or more normalized dimension tables that surround a centralized fact table. The centralized fact table contains data elements of interest that are derived from the dimension tables.

For more information on the STARJOIN facility, see the section, "SAS Scalable Performance Data (SPD) Server STARJOIN Facility," in the chapter "SAS Scalable Performance Data (SPD) Server," and the chapter, "SAS Scalable Performance Data (SPD) Server STARJOIN Facility," in the *SAS Scalable Performance Data Server: User's Guide*.

SAS Scalable Performance Data (SPD) Server 4.3 Index Scans

SAS Scalable Performance Data (SPD) Server 4.3 SQL enables users to perform fast index scans on large tables. Rather than scanning entire tables that might have millions of rows, in specific cases, SAS Scalable Performance Data (SPD) Server 4.3 SQL can use index data to resolve the query. Index data is compact, small, and faster to query than an entire table. SAS Scalable Performance Data (SPD) Server 4.3 SQL provides enhanced index scan support for the following functions:

MIN, MAX, COUNT, COUNT DISTINCT, NMISS, RANGE

For more information on index scans, see the section, "SAS Scalable Performance Data (SPD) Server Index Scan," in the chapter "SAS Scalable Performance Data (SPD) Server SQL Features," in the *SAS Scalable Performance Data Server: User's Guide*.

SAS Scalable Performance Data (SPD) Server 4.3 Optimized Correlated Queries

Intelligent storage must have the ability to interpret and process complex requests such as correlated queries. A correlated query is a select expression where a predicate within the query has a relationship to a column that is defined in another layer of code. Today's business and analytic intelligence tools often generate SQL queries that are nested three or four layers deep. Queries with cross-nested relationships use significant processor

resources and require more time to complete processing. New algorithms in the SQL planner of SAS Scalable Performance Data (SPD) Server 4.3 implement techniques that significantly improve the performance of correlated queries for patterns that permit query rewrites or query de-correlation.

SAS Scalable Performance Data (SPD) Server 4.3 **Parallel GROUP BY**

Parallel GROUP BY is a high performance parallel summarization of data that is executed using SQL. Parallel GROUP BY works against single tables that are used to aggregate data. Summarization tasks are common in data warehousing applications. Parallel GROUP BY was developed to quicken processor performance summarization tasks. Parallel GROUP BY is often used in SQL queries (through the use of sub-queries) to apply selection lists for inclusion or exclusion.

Parallel GROUP BY support in SAS Scalable Performance Data (SPD) Server 4.3 has been expanded. Parallel GROUP BY is integrated in the WHERE clause planner code so that it will boost the capabilities of the SAS Scalable Performance Data (SPD) Server SQL engine. Any section of code that matches the parallel GROUP BY trigger pattern will use parallel GROUP BY.

For more information on parallel GROUP BY, see the section, "Parallel Joins with GROUP BY," in the chapter "SAS Scalable Performance Data (SPD) Server SQL Features," in the *SAS Scalable Performance Data Server: User's Guide*.

SAS Scalable Performance Data (SPD) Server 4.3 **Parallel Join**

Parallel join is a high-performance pairwise join of two SAS Scalable Performance Data (SPD) Server tables. The parallel join feature enhances join performance in two ways. First, SAS Scalable Performance Data (SPD) Server parallel joins are performed using parallel threading. Second, SAS Scalable Performance Data (SPD) Server parallel joins use enhanced data summarization methods after rows in a table are joined.

For more information on SAS Scalable Performance Data (SPD) Server parallel joins, see the section, "Parallel Join Facility," in the chapter "SAS Scalable Performance Data (SPD) Server SQL Features," in the *SAS Scalable Performance Data Server: User's Guide*.

SAS Scalable Performance Data (SPD) Server 4.3

MINMAX Table Indexing

SAS Scalable Performance Data (SPD) Server 4.3 contains a new table option called MINMAXVARLIST=. The primary purpose of the MINMAXVARLIST= table option is for use with SAS Scalable Performance Data (SPD) Server 4.3 dynamic cluster tables, where specific member tables in the dynamic cluster contain a set or range of values, such as sales data for a given month. When an SAS Scalable Performance Data (SPD) Server SQL subsetting WHERE clause specifies specific months from a range of sales data values, the WHERE clause planner checks the MINMAX indexes. Based on the MINMAX index information, the SAS Scalable Performance Data (SPD) Server WHERE clause planner includes or eliminates member tables in the dynamic cluster for evaluation.

Use the MINMAXVARLIST= table option with numeric columns. MINMAXVARLIST= uses the list of columns you submit to build an index. The MINMAXVARLIST= index contains only the minimum and maximum values for each column. The WHERE clause planner uses the index to filter SQL predicates quickly, and to include or eliminate member tables in the dynamic cluster for evaluation.

Although the MINMAXVARLIST= table option is primarily intended for use with dynamic cluster tables, it can be used with standard SAS Scalable Performance Data (SPD) Server tables. MINMAXVARLIST= can help reduce the need to create many indexes on a table, which can save valuable resources and space.

For more information on SAS Scalable Performance Data (SPD) Server the MINMAXVARLIST= table options, see the section, "MINMAX Indexes," in the chapter, "Optimizing SAS Scalable Performance Data (SPD) Server Performance," in the *SAS Scalable Performance Data Server: User's Guide*.

SAS Scalable Performance Data (SPD) Server 4.3 WHERE Costing Improvements

The WHERE clause planner that is implemented in SAS Scalable Performance Data (SPD) Server 4.3 avoids resource-intensive computations and uses simple computations where possible. WHERE clauses in large database operations can be resource-intensive. In SAS Scalable Performance Data (SPD) Server 3.x and earlier releases, users often needed to manually tune queries for performance. Tuning was accomplished using macro variables and index settings. The WHERE clause planner that is integrated into SAS Scalable Performance Data (SPD) Server 4.3 does the tuning for the user by costing the different approaches to index evaluation.

For more information on SAS Scalable Performance Data (SPD) Server WHERE clause improvements, see the section, "Optimizing WHERE Clauses," in the chapter, "Optimizing SAS Scalable Performance Data Server Performance," in the *SAS Scalable Performance Data Server: User's Guide*.

SAS Scalable Performance Data (SPD) Server 4.3 Cluster Tables

SAS Scalable Performance Data (SPD) Server 4.3 uses a virtual table structure called a cluster table. Cluster tables provide a storage architecture that has parallel processing and data management capabilities.

A cluster table is a structure that can store multiple SAS Scalable Performance Data (SPD) Server tables. A cluster table is composed of member tables (or partitions). Each member can store a single SAS Scalable Performance Data (SPD) Server table. The cluster table uses a layer of metadata to manage the members. Cluster tables can also be used in WHERE clause costing. Each member in a cluster table is analyzed and assigned an EVAL strategy that best fits the data patterns in the member or slot. Using multiple EVAL strategies while performing WHERE clause costing on a cluster table provides better process granularity, which can improve overall data throughput and performance.

Dynamic Cluster Tables

SAS Scalable Performance Data (SPD) Server cluster tables are virtual table structures. SAS Scalable Performance Data (SPD) Server 4.3 cluster tables are a bound collection of multiple members. Each member is a standard SAS Scalable Performance Data (SPD) Server table. All member tables that belong to a dynamic cluster table must share the same metadata formats and organization. SAS Scalable Performance Data (SPD) Server 4.3 dynamic cluster tables use metadata to manage the data that is contained in the member tables.

The SAS Scalable Performance Data (SPD) Server 4.3 dynamic cluster table structure provides architecture that enables flexible loading, rapid storage, and parallel processing for very large data tables. Using dynamic cluster tables, loading data, removing data, and refreshing tables in very large data marts become easier and more timely. Dynamic cluster tables provide organizational features and performance benefits that traditional SAS tables and SAS Scalable Performance Data (SPD) Server tables do not have.

For example, you can add new data or remove historical data from very large tables by accessing only the member tables that are affected by the change. You can access the individual member tables in parallel. This strategy reduces the time that is needed for the job to complete and uses very simple commands. Furthermore, a complete refresh of a dynamic cluster table can occur using a fraction of the disk space that is needed to refresh a large traditional SAS or SAS Scalable Performance Data (SPD) Server table that contains the same amount of data.

For more information on SAS Scalable Performance Data (SPD) Server dynamic cluster tables, see the chapter, "SAS Scalable Performance Data (SPD) Server Dynamic Cluster Tables," in the *SAS Scalable Performance*

Unsupported Features in Cluster Tables

Because of differences in the load and read structures for dynamic cluster tables, some standard features that are normally available in Base SAS tables and SAS Scalable Performance Data (SPD) Server tables are currently not supported in SAS Scalable Performance Data (SPD) Server 4.3 cluster tables.

The features are:

- You cannot append data to a dynamic cluster table. To append data to a dynamic cluster table, the table must be unclustered, the data is appended to the individual unclustered files, and then the unclustered files must be reclustered.
- Record-level locking is not allowed.
- The SAS Scalable Performance Data (SPD) Server backup/restore utility is not available.
- Copying data with PROC COPY or PROC SQL is not supported.

If a task for a dynamic cluster table requires one of these features, you should undo the dynamic cluster table and create standard SAS Scalable Performance Data (SPD) Server tables.

SAS Scalable Performance Data (SPD) Server 4.3 Random Placement of Initial Data Partition Files in DATAPATH= Setting

In SAS Scalable Performance Data (SPD) Server 3.x, the initial data partition files for all tables in the same domain are assigned to the first DATAPATH= setting that was defined in the libnames.parm LIBNAME configuration file. Subsequent data partition files for a table are placed in subsequent DATAPATHs. When all SAS Scalable Performance Data (SPD) Server DATAPATHs contain a data partition file, the process returns to the first DATAPATH and continues. However, numerous SAS Scalable Performance Data (SPD) Server installations have many small-to-medium-sized tables that do not have data partition files in all of the available DATAPATHs. This process could unevenly balance the distribution of data on the disk, resulting with the first few DATAPATHs in a domain containing significantly more data than the last few DATAPATHs in the domain. The uneven data distribution results in unbalanced I/O.

In SAS Scalable Performance Data (SPD) Server 4.3, the initial data partition files for all

tables in the same domain are no longer assigned to the first DATAPATH= setting that was defined in the libnames.parm LIBNAME configuration file. Instead, SAS Scalable Performance Data (SPD) Server randomly chooses from the available DATAPATHs when assigning the initial data partition files for a large data table. As a result, the data is distributed more evenly and permits more balanced I/O within SAS Scalable Performance Data (SPD) Server processing.

By default, SAS Scalable Performance Data (SPD) Server 4.3 is configured to use random placement of initial data partition files among SAS Scalable Performance Data (SPD) Server DATAPATHs. The RANDOMPLACEDPF option is specified in the spdsserv.parm file. To disable random placement of initial data partition files in the DATAPATH= list, remove the RANDOMPLACEDPF option from your spdsserv.parm file.

For more information on the RANDOMPLACEDPF option, see the chapter, "Setting up SAS Scalable Performance Data (SPD) Server Parameter Files," in the *SAS Scalable Performance Data Server: Administrator's Guide*.

SAS Scalable Performance Data (SPD) Server 4.3 **Debugging Tools**

SAS Scalable Performance Data (SPD) Server 4.3 includes useful debugging tools. The debugging tools enable SAS Scalable Performance Data (SPD) Server system administrators to create debug images and to evaluate test images without interfering with a production SAS Scalable Performance Data (SPD) Server environment. The debugging tools are for use with SAS Scalable Performance Data (SPD) Server 4.3 running on SAS 9.1.3. The debugging tools are organized into LIBNAME statement options for debugging, and server parameter file options for debugging.

For more information on SAS Scalable Performance Data (SPD) Server debugging tools, see the chapter, "SAS Scalable Performance Data (SPD) Server Debugging Tools," in the *SAS Scalable Performance Data Server: Administrator's Guide*.

Copyright © 2008 by SAS Institute Inc., Cary, NC, USA. All rights reserved.

SAS Scalable Performance Data (SPD) Server Overview

- [Introduction to SAS Scalable Performance Data \(SPD\) Server](#)
 - [The SAS Scalable Performance Data \(SPD\) Server Client/Server Model](#)
 - [Symmetric Multi-Processor Hosts](#)
 - [SAS Scalable Performance Data \(SPD\) Server Host Services for Clients](#)
 - [Accessing SAS Scalable Performance Data \(SPD\) Server Using SAS](#)
 - [Securing SAS Data](#)
 - [Organizing SAS Data](#)
 - [SAS Scalable Performance Data \(SPD\) Server Performance Enhancements](#)
 - [SAS Scalable Performance Data \(SPD\) Server Extensions to Base SAS](#)
 - [Using SAS Scalable Performance Data \(SPD\) Server with Data Warehousing](#)
-

[Introduction to SAS Scalable Performance Data \(SPD\) Server](#)

SAS Scalable Performance Data (SPD) Server software is designed for high-performance data delivery. Its primary function is to provide user access to SAS data for intensive processing (queries and sorts) on the host server machine. When client workstations from varying operating platforms send processing requests to a SAS Scalable Performance Data (SPD) Server host, the host returns results in the format required by each client workstation. SAS Scalable Performance Data (SPD) Server uses the power of parallel processing to exploit the threading capabilities of servers with multiple processors.

SAS Scalable Performance Data (SPD) Server executes *threads*, units of processing, in parallel on a SAS Scalable Performance Data (SPD) Server host. The software tasks are performed in conjunction with an operating system that enables threads to execute on any of the machine's available processors. A specialized machine and operating system are important processing partners, but SAS Scalable Performance Data (SPD) Server's power is derived from the software architecture that enable it to rapidly and efficiently process SAS data in concurrent parallel threads on multiple processors .

SAS Scalable Performance Data (SPD) Server is the high-speed processing tool among SAS products. SPD 4.3 introduces on-disk structures that are compatible with SAS 9 and the large table capacities that it supports. Enterprise-wide data mining often creates immense tables. In order to generate business intelligence quickly, the ability to update tables that contain billions of rows is more important than ever. The cluster table structure introduced in SAS Scalable Performance Data (SPD) Server 4.3 provides a new foundation for the next generation of SAS data storage. Previous versions of SAS Scalable Performance Data (SPD) Server were based on 32-bit architecture that supported just over 2 billion rows and 32,768 columns. SAS Scalable Performance Data (SPD) Server is based on a 64-bit architecture which supports tables with over 9 *quintillion* rows and over 2 *billion columns*.

SAS Scalable Performance Data (SPD) Server 4.4 operates on computers running SAS 9.1.3 or later. PC users that do not use SAS can still use SAS Scalable Performance Data (SPD) Server. Information on connecting to SAS Scalable Performance Data (SPD) Server with Other Clients is found in [Using SAS Scalable Performance Data \(SPD\) Server With Other Clients](#). SAS users can access SAS Scalable Performance Data (SPD) Server either by using SQL pass-through or by using SAS language.

Syntax Conventions: SAS Scalable Performance Data (SPD) Server software supports both SAS users and non-SAS users. The SAS Scalable Performance Data (SPD) Server document uses common terminology that both audiences should understand. In the SAS Scalable Performance Data (SPD) Server documentation, SAS data sets are referred to as tables, SAS variables are referred to as columns, and SAS observations are referred to as rows. The SAS Scalable Performance Data (SPD) Server product is referred to as SAS Scalable Performance Data (SPD) Server or "the software", depending on the context of the documentation.

The SAS Scalable Performance Data (SPD) Server Client/Server Model

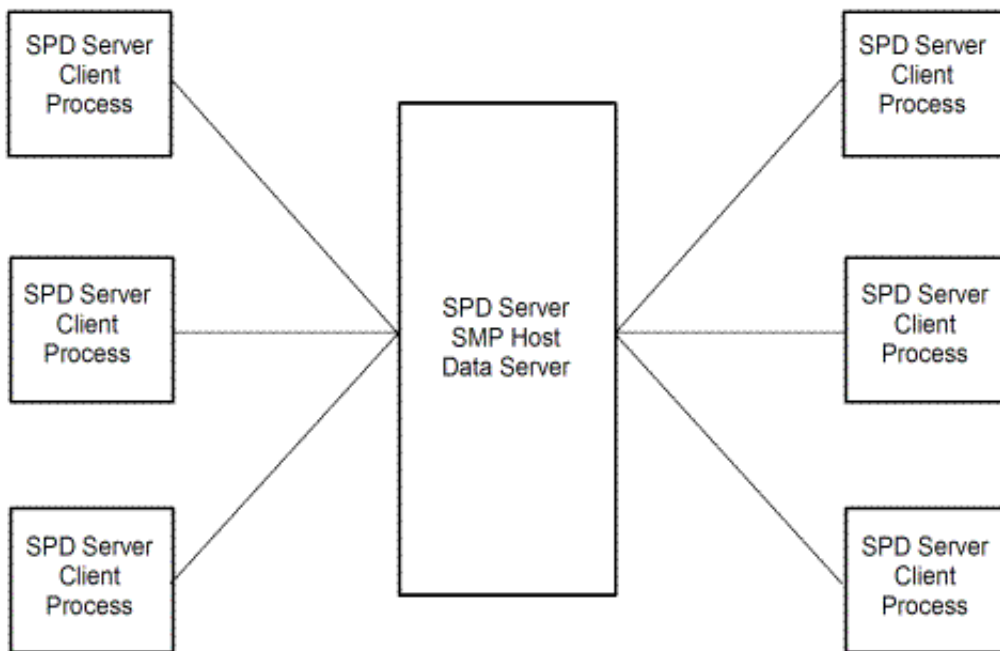
SAS Scalable Performance Data (SPD) Server software divides SAS processing loads between the client and server. The [Client/Server Model](#) diagram shows a simple client/server topology. The server hosts multiple concurrent clients while performing the heaviest processing tasks. Typical clients are desktop PCs or low-end UNIX workstations running front-end software. The front-end application sends the client's data requests over the network to the server and processes the information that the server returns.

You can create one or more SAS Scalable Performance Data (SPD) Servers on the host server machine. When a SAS Scalable Performance Data (SPD) Server host receives a client's data request, it performs some action on behalf of the client. The action varies with the request received.

Where does the user fit within in the SAS Scalable Performance Data (SPD) Server Client/Server model? Users initiate SAS Scalable Performance Data (SPD) Server client sessions. In this documentation, the term 'user' refers to the operator of a SAS Scalable Performance Data (SPD) Server client.

[Figure 1.1](#)

The SAS Scalable Performance Data (SPD) Server Client/Server Model



Symmetric Multi-Processor Hosts

SAS Scalable Performance Data (SPD) Server host machines use operating systems that can process concurrent threads in

parallel on multiple processors. SAS Scalable Performance Data (SPD) Server exploits symmetric multiprocessing (SMP) hardware and software architecture.

The number of processors on an SMP server varies by manufacturer and model. The operating system of the machine must also support the parallel processing. Operating systems which possess a threaded kernel enjoy enhanced performance because the threaded kernel prevents contention issues among competing threads in real-time. Synergy between processors and threads allows SAS Scalable Performance Data (SPD) Server to scale processing performance. The scalability, in turn, significantly improves the speed of SAS Scalable Performance Data (SPD) Server table creates, appends, scans, queries, and sorts.

SAS Scalable Performance Data (SPD) Server Host Services for Clients

SAS Scalable Performance Data (SPD) Server hosts provide multiple services to SAS Scalable Performance Data (SPD) Server clients:

- **Access to data stores** SAS Scalable Performance Data (SPD) Server offers concurrent read access and retrieval of SAS data.
- **High-speed data server** SAS Scalable Performance Data (SPD) Server manages and processes massive SAS tables.
- **Offloads heavy processing work** SAS Scalable Performance Data (SPD) Server divides the labor. The Server process retrieves, sorts, and subsets SAS data. A client process reviews and analyzes the data that the Server returns.
- **Embellishes client hardware** SAS Scalable Performance Data (SPD) Server host machines are able to utilize the computing hardware resources that are required to process large tables efficiently and rapidly.
- **Reduces network traffic** SAS Scalable Performance Data (SPD) Servers read, sort, and subset entire SAS tables and then return answer sets. A query subset replaces large file downloads to the client machine. SAS Scalable Performance Data (SPD) Server also offers a common storage facility. Multiple client users can use the same SAS data on the server without having to each transfer the SAS data to their workstations.
- **Provides multi-platform support** SAS Scalable Performance Data (SPD) Server allows clients to share SAS data across computing platforms with other SAS users.

Table 1. 1

SAS Scalable Performance Data (SPD) Server Features

SPD Server Feature	SPD Server Client Action	SPD Server Host Response
Support for Gigabytes of data	The SPD Server client inputs existing SAS tables with a PROC COPY statement or creates a SPD Server table using a SAS data step or procedure. SPD Server clients can also use SQL pass-through CREATE, COPY, or LOAD statements to input SAS tables.	The SPD Server host creates component files that are composed of one or more physical partition files. The server stores the physical partition files in one or more device / directory paths.

Scalable Symmetric Multiple Processor (SMP) Support	The SPD Server client runs SAS procedures and SQL pass-through syntax to read, sort, index, or query an SPD Server table.	The SPD Server host utilizes its threaded operating system to perform concurrent processing tasks distributed across multiple processors.
Selective Parallel Queries	The SPD Server client uses WHERE-clause or SQL SELECT syntax. Pass-through SQL, PROC SQL, and non-SAS WHERE alternatives are supported.	The SPD Server host supports and subsets SPD Server tables, then delivers query answer sets to clients.
Parallel Loads	The SPD Server client runs SAS procedures with LOAD or COPY to store SAS data and indexes.	The SPD Server host uses multiple threads to load and store tables and indexes.
Parallel Indexes	The SPD Server client creates table indexes using a DATA step or the DATASETS procedure with an INDEX option, or pass-through SQL with the LOAD or COPY command.	The SPD Server host creates SPD Server table indexes in parallel.
SAS Data Security	The SPD Server client accesses the SPD Server host using SQL pass-through, a LIBNAME statement, or a non-SAS alternative connection.	The SPD Server host secures SPD Server files at the LIBNAME domain and / or table, column, and row level.

[Accessing SAS Scalable Performance Data \(SPD\) Server Using SAS](#)

You begin a SAS Scalable Performance Data (SPD) Server session by starting your SAS Scalable Performance Data (SPD) Server client. There are two ways to start your SAS Scalable Performance Data (SPD) Server client session. You can use SQL pass-through commands to start your SAS Scalable Performance Data (SPD) Server client session, or you can use a LIBNAME statement to start your SAS Scalable Performance Data (SPD) Server client session. Both methods use the SASSPDS engine and initiate communication between the SAS Scalable Performance Data (SPD) Server client machine and SAS Scalable Performance Data (SPD) Server host.

- [SQL Pass-Through Facility](#)
- [LIBNAME Access](#)
- [SAS Scalable Performance Data \(SPD\) Server Host Name Server](#)
- [Specifying the Port Address for the Name Server](#)

[SQL Pass-Through Facility](#)

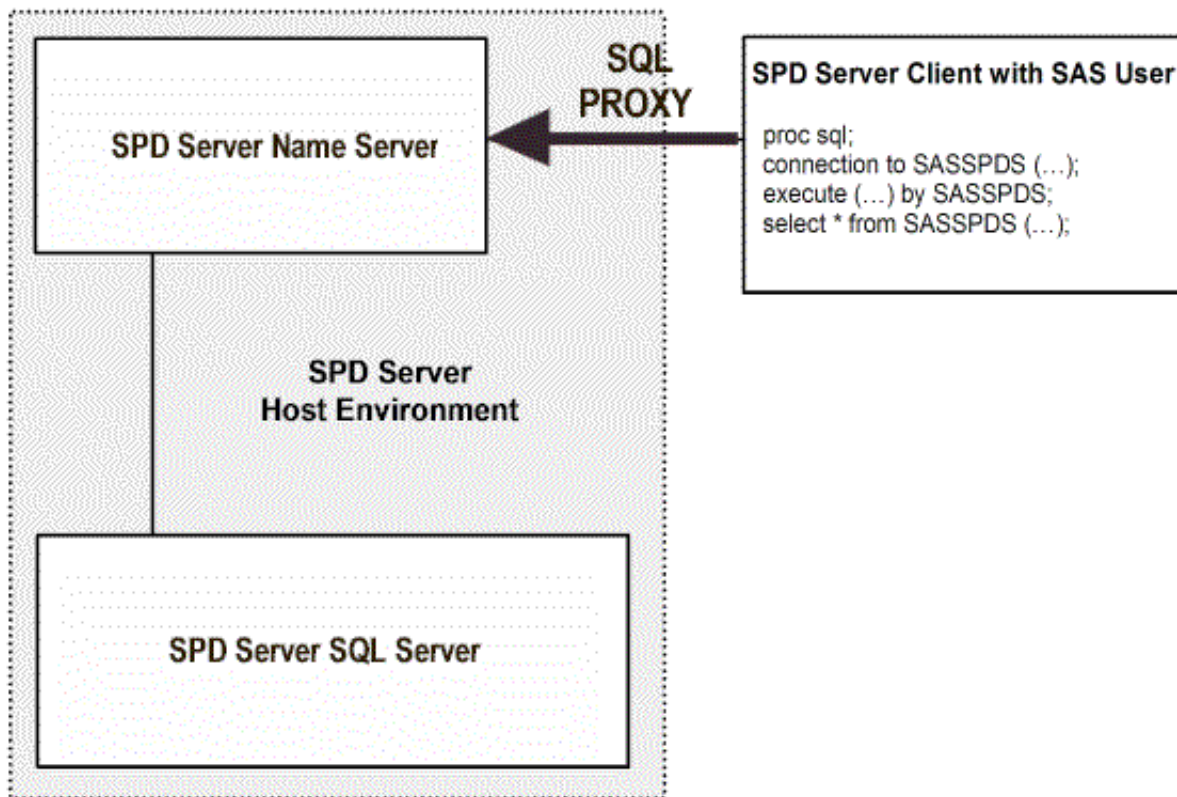
SAS Scalable Performance Data (SPD) Server can use SQL pass-through commands. The SAS Scalable Performance Data (SPD) Server host can perform complete SQL-expression evaluation. SAS Scalable Performance Data (SPD) Server also supports nested SQL pass-through commands. Nested SQL pass-through commands permit you to connect to other SAS Scalable Performance Data (SPD) Server hosts while you are still

connected to your SAS Scalable Performance Data (SPD) Server host. You can use nested pass-through commands to distribute simultaneous SQL queries across multiple SAS Scalable Performance Data (SPD) Server hosts on your network.

The SQL pass-through facility can be accessed with or without SAS syntax and applications. You can use SAS to connect to an SAS Scalable Performance Data (SPD) Server host by using pass-through syntax from PROC SQL or from other SQL-aware SAS applications. The chapter on [Accessing and Creating SAS Scalable Performance Data Server Tables](#) contains more detailed information about the SAS Scalable Performance Data (SPD) Server pass-through facility and provides examples of the syntax.

Figure 1.2

SAS Scalable Performance Data (SPD) Server Client Access to SAS Scalable Performance Data (SPD) Server Host Using SQL Pass-Through and SAS CONNECT

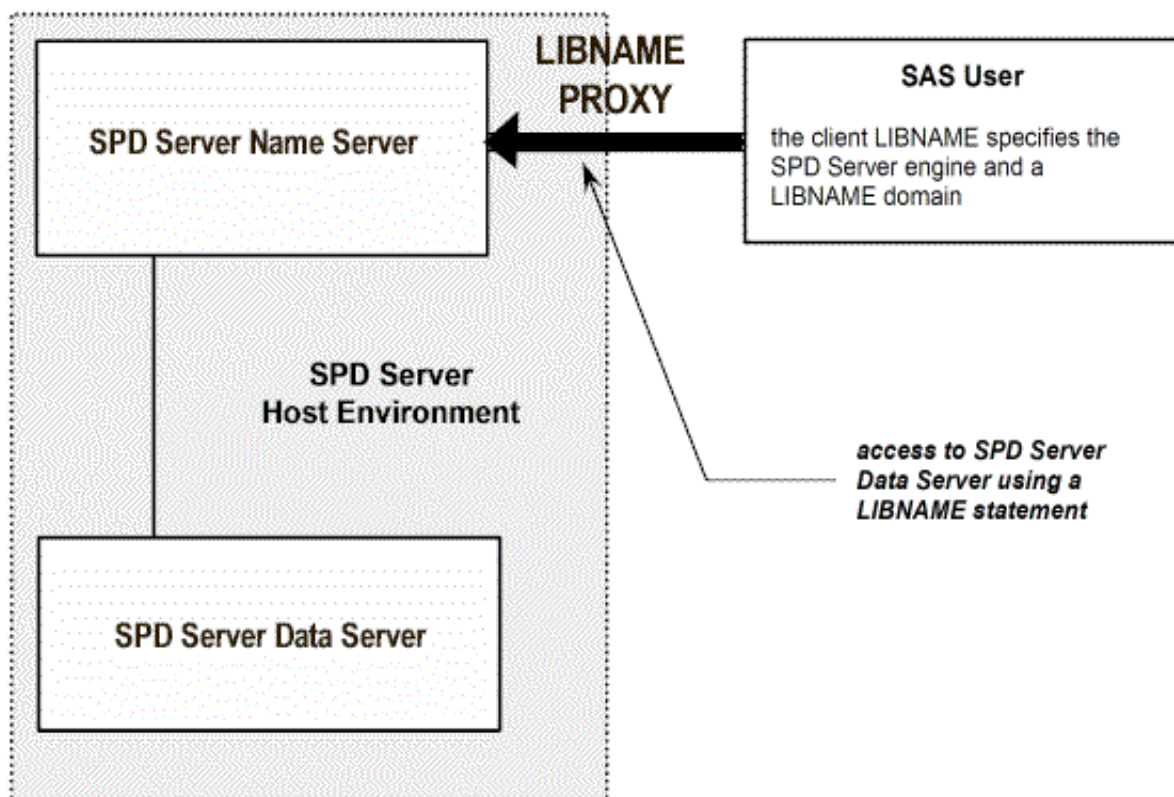


LIBNAME Access

SAS users can initiate a client session by issuing a LIBNAME statement using the engine SASSPDS. LIBNAME access is illustrated in Figure 1.3. The documentation chapter on [Connecting SAS Clients to SAS Scalable Performance Data \(SPD\) Server](#) explains the mechanics of LIBNAME access to the engine and SAS Scalable Performance Data (SPD) Server LIBNAME options.

Figure 1.3

SAS Scalable Performance Data (SPD) Server Client (SAS User) Access to SAS Scalable Performance Data (SPD) Server Host Using a LIBNAME Statement



SAS Scalable Performance Data (SPD) Server Host Name Server

Distributed computing may enrich user resources, but it has an inherent problem. To connect to a SAS Scalable Performance Data (SPD) Server, you must know its location within your network. Instead of requiring users to memorize long paths or IP addresses, SAS Scalable Performance Data (SPD) Server software uses a specialized server called a *name server*. The SAS Scalable Performance Data (SPD) Server name server locates active SAS Scalable Performance Data (SPD) Server hosts on your network. A name server recognizes active SAS Scalable Performance Data (SPD) Server machines because all the SAS Scalable Performance Data (SPD) Servers 'register' with the name server as they come up and contact the host machine.

The name server keeps network addresses and a list of the LIBNAME domains for each SAS Scalable Performance Data (SPD) Server host. What is an SAS Scalable Performance Data (SPD) Server LIBNAME domain? A SAS Scalable Performance Data (SPD) Server LIBNAME domain is a logical entity that SAS Scalable Performance Data (SPD) Server creates. A LIBNAME domain maintains domain attributes such as the library name, owner, and contents. Whenever you use a LIBNAME statement to specify a LIBNAME domain, a name server can determine the correct directory path to the SAS Scalable Performance Data (SPD) Server data library and connect your SAS Scalable Performance Data (SPD) Server client to the SAS Scalable Performance Data (SPD) Server host for that domain.

Specifying the Port Address for the Name Server

SAS Scalable Performance Data (SPD) Server clients use port addressing to locate SPD name servers. SAS Scalable Performance Data (SPD) Server administrators must assign a port address to a name server. Most UNIX system clients use their local `/etc/services` file to register port assignments. The service name for a SAS Scalable Performance Data (SPD) Server name server in an `/etc/services` file must be **SPDSNAME**. PC clients use services files to register port assignments. The services files on PC clients vary according to the software that the PC network uses.

When a client SAS Scalable Performance Data (SPD) Server application issues a LIBNAME statement that does not contain the port address of the name server, SAS Scalable Performance Data (SPD) Server checks the services file for the SPDSNAME entry and the port address. Registering the name server port assignment in your client's network services file relieves you from the responsibility of coding name server port numbers when you write SAS jobs. The Help on Connecting SAS Clients to SAS Scalable Performance Data (SPD) Server contains examples that show you how to [Connect to SAS Scalable Performance Data \(SPD\) Server Using a LIBNAME Statement and a Name Server](#).

Securing SAS Data

- [LIBNAME Domain Registration](#)
- [ACL File Security](#)

LIBNAME Domain Registration

The name server helps SAS Scalable Performance Data (SPD) Server clients locate and connect to SAS Scalable Performance Data (SPD) Server hosts. The name server also controls access to the SAS Scalable Performance Data (SPD) Server LIBNAME domains. How does the name server get domain information? The SAS Scalable Performance Data (SPD) Server administrator defines LIBNAME domains in an SAS Scalable Performance Data (SPD) Server LIBNAME parameter file.

When a SAS Scalable Performance Data (SPD) Server administrator brings up a server on the host machine, SAS Scalable Performance Data (SPD) Server reads the `spdsrv.parm` parameter file and registers the domains that are listed in the parameter file with the name server. The name server remembers which SAS Scalable Performance Data (SPD) Server host or hosts have access to a given LIBNAME domain. If you want to specify a LIBNAME domain, you can do so using a LIBNAME statement or a pass-through SQL CONNECT statement. Your SAS Scalable Performance Data (SPD) Server administrator can provide you with a list of the LIBNAME domains that are mapped to your SAS Scalable Performance Data (SPD) Server host machine.

ACL File Security

SAS Scalable Performance Data (SPD) Server uses Access Control Lists (ACLs) and SAS Scalable Performance Data (SPD) Server user IDs to secure domain resources. You obtain your user ID and password from your SAS Scalable Performance Data (SPD) Server administrator.

SAS Scalable Performance Data (SPD) Server also supports ACL groups, which are similar to UNIX groups. SAS

Scalable Performance Data (SPD) Server administrators can associate a SAS Scalable Performance Data (SPD) Server user as many as five ACL groups.

ACL file security is turned on by default when an administrator brings up SAS Scalable Performance Data (SPD) Server. ACL permissions affect all SAS Scalable Performance Data (SPD) Server resources, including domains, tables, table columns, catalogs, catalog entries, and utility files. When ACL file security is enabled, SAS Scalable Performance Data (SPD) Server only grants access rights to the owner (creator) of a SAS Scalable Performance Data (SPD) Server resource. Resource owners can use PROC SPDO to grant ACL permissions to a specific group (called an ACL group) or to all SAS Scalable Performance Data (SPD) Server users.

The resource owner can use the following properties to grant ACL permissions to all SAS Scalable Performance Data (SPD) Server users:

READ

universal READ access to the resource (read or query).

WRITE

universal WRITE access to the resource (append to or update).

ALTER

universal ALTER access to the resource (rename, delete, or replace a resource and add, delete indexes associated with a table).

The resource owner can use the following properties to grant ACL permissions to a named ACL group:

GROUPREAD

group READ access to the resource (read or query).

GROUPWRITE

group WRITE access to the resource (append to or update).

GROUPALTER

group ALTER access to the resource (rename, delete, or replace a resource and add, delete indexes associated with a table).

[Organizing SAS Data](#)

- [SAS Scalable Performance Data \(SPD\) Server Tables](#)
- [SAS Scalable Performance Data \(SPD\) Server Component Files](#)
- [SAS Scalable Performance Data \(SPD\) Server Table Indexes](#)

[SAS Scalable Performance Data \(SPD\) Server Tables](#)

SAS Scalable Performance Data (SPD) Server software alters SAS tables to enable high-performance processing. SAS Scalable Performance Data (SPD) Server tables are physically different than a Base SAS table. You can use tables in either SAS or native SAS Scalable Performance Data (SPD) Server format. The SAS Scalable Performance Data (SPD) Server User's Guide chapter on [Accessing and Creating SAS Scalable Performance Data \(SPD\) Server Tables](#) discusses how a simple SAS PROC COPY statement handles [conversion details and](#)

[changing between table formats.](#)

How are SAS tables organized? SAS tables stores a single file that contains the data descriptors and the table data. The data are column values, the descriptors are metadata that describe the column and data formatting that the table uses.

SAS Scalable Performance Data (SPD) Server tables do not reuse space. When an SQL command to delete one or more rows from a table is issued, the row is marked deleted and the space will not be reused. To recapture the space, the table must be copied.

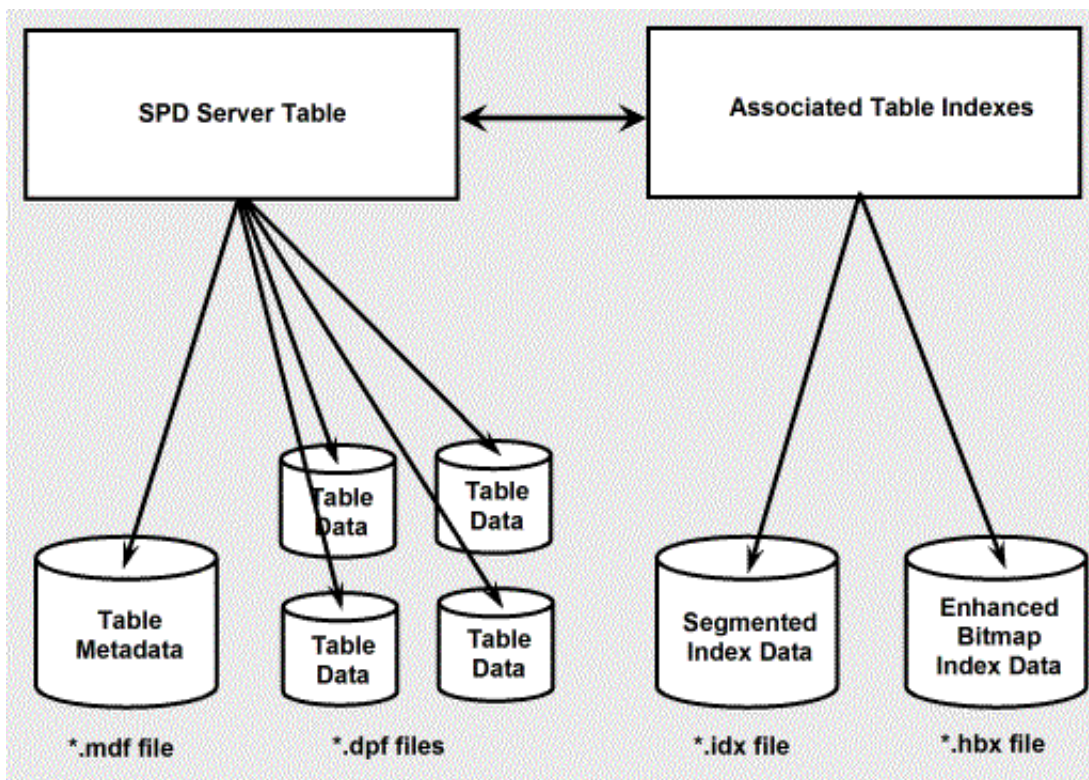
The diagram of the [SAS Scalable Performance Data \(SPD\) Server Table Component Files](#) shows differences in the architecture between SAS Scalable Performance Data (SPD) Server tables and SAS tables. SAS Scalable Performance Data (SPD) Server uses component files to store tables. One component file stores the stream of data values. Another component file stores the column and data descriptors, the metadata. If you create an index for a column or a composite of columns, SAS Scalable Performance Data (SPD) Server creates component files for each index.

SAS Scalable Performance Data (SPD) Server Component Files

SAS Scalable Performance Data (SPD) Server uses four types of component files to store SAS Scalable Performance Data (SPD) Server tables. The diagram of the [SAS Scalable Performance Data \(SPD\) Server Table Component Files](#) shows the components of SAS Scalable Performance Data (SPD) Server tables. Two component files store table information: the *.dpf component file stores a stream of the table's data values, and the *.mdf component file stores the table's metadata (column and data descriptors) information. SAS Scalable Performance Data (SPD) Server also creates two more component files to manage index data: *.hbx components are unique global B-tree indexes and *.idx components are segmented views of the indexed column data. The *.idx components are especially useful in evaluating parallel WHERE-clauses.

[Figure 1.4](#)

SAS Scalable Performance Data (SPD) Server Component Files



SAS Scalable Performance Data (SPD) Server partitions component files when they are created to keep them from growing too large. Each partitioned component file is stored as one or more disk files. There are several advantages to partitioning the component files:

- Very Large Tables:** SAS Scalable Performance Data (SPD) Server bypasses file size limits imposed by many applications and operating systems. By using partitioned component files, SAS Scalable Performance Data (SPD) Server can support any file system transparently.
- Multiple Directory Paths:** SAS Scalable Performance Data (SPD) Server can access data libraries that span numerous directory paths and storage devices. SAS Scalable Performance Data (SPD) Server software partitions massive data libraries into component files. The component architecture enables rapid threaded data access while circumventing device capacity and file size limitation issues. Storage lists transparently track component file locations so users can access multiple storage devices as a single volume, even if file partitions exist in different locations.
- Flexibility in Storage:** There is no need to store data tables and associated indexes in the same location when using SAS Scalable Performance Data (SPD) Server component files. Data files and associated indexes can be stored on different directory structures or devices if you wish. When deciding where to store component SAS Scalable Performance Data (SPD) Server tables, you only need to consider the cost, performance, and availability of the disk space.
- Improved Table Scan Performance:** Data component partitions that are created using fixed-size intervals will perform aggressively during parallelized full table scans. The documentation chapter on [SAS Scalable Performance Data \(SPD\) Server Table Options](#) contains information on how to use the `PARTSIZE=` option to control partition size.

SAS Scalable Performance Data (SPD) Server Table Indexes

SAS Scalable Performance Data (SPD) Server allows you to create indexes on table columns. SAS Scalable Performance Data (SPD) Server can thread WHERE-clause evaluations for tables that are not indexed. Indexes enable more rapid WHERE-clause evaluations. Large tables in particular should be indexed to exploit SAS Scalable Performance Data (SPD) Server performance. A detailed description of the SAS Scalable Performance Data (SPD) Server index is provided in the Usage section on [Indexing a Table](#).

SAS Scalable Performance Data (SPD) Server Performance Enhancements

- [SAS Scalable Performance Data \(SPD\) Server Pass-Through SQL Enhancements](#)
- [Implicit and Explicit Server Sorts](#)
- [Modified SAS Heapsort](#)
- [Indexed Parallel Table Scan](#)
- [Improved Table Appends](#)

SAS Scalable Performance Data (SPD) Server Pass-Through SQL Enhancements

You can use pass-through SQL to submit SQL statements that use SAS Scalable Performance Data (SPD) Server tables directly to SAS Scalable Performance Data (SPD) Server. The SAS Scalable Performance Data (SPD) Server SQL planner contains several optimizations that you can utilize to create SQL queries that can take advantage of symmetric multiprocessing and SPD table indexes, resulting in improved SQL query performance. Refer to the SAS Scalable Performance Data (SPD) Server User's Guide section on the [SAS Scalable Performance Data \(SPD\) Server SQL Planner](#) for more information on SAS Scalable Performance Data (SPD) Server pass-through SQL enhancements.

Implicit and Explicit Server Sorts

You can use implicit or explicit sorts with SAS Scalable Performance Data (SPD) Server. For example, the PROC SORT in Base SAS software is an explicit sort. You can use PROC SORT with SAS Scalable Performance Data (SPD) Server as well.

An implicit sort is unique to SAS Scalable Performance Data (SPD) Server. Each time you submit a SAS statement with a BY clause, SAS Scalable Performance Data (SPD) Server sorts your data -- unless the table is already sorted or indexed on the BY column. The automatic sort is very convenient. The documentation chapter on [Accessing and Creating SAS Scalable Performance Data Server Tables](#) contains tips on how and when to use each sort type.

Modified SAS Heapsort

SAS Scalable Performance Data (SPD) Server uses Heapsort as its default sort with some slight changes. Under SAS Scalable Performance Data (SPD) Server, Heapsort compares available memory on the server to the memory required to load and process the index key data in memory. If the memory is not constrained, SAS Scalable Performance Data (SPD) Server performs the Heapsort in RAM memory.

Indexed Parallel Table Scan

SAS Scalable Performance Data (SPD) Server indexes are designed to support parallelism. Experienced RDBMS users are accustomed to a perceptible processing lag that occurs when databases must read or process enormous tables. When SAS Scalable Performance Data (SPD) Server performs table queries, the SAS Scalable Performance Data (SPD) Server index architecture enables the software to analyze different table sections or segments in parallel. By processing large table segments in parallel, SAS Scalable Performance Data (SPD) Server delivers much faster data throughput. The faster throughput may be difficult to perceive on small tables, but when SAS Scalable Performance Data (SPD) Server performs scans on very large tables, the processing performance is significantly faster than database systems that support only serial indexed table scans.

Improved Table Appends

SAS Scalable Performance Data (SPD) Server decomposes table append operation into a set of steps that can be performed in parallel. The level of parallelism attained depends on the number of indexes present on the table. The more indexes you have, the greater the potential exploitation of parallelism during the append processing.

Tip: You can save time by creating an empty table in SAS Scalable Performance Data (SPD) Server, then define your indexes on it, and then append the data, as opposed to loading the table and then creating the indexes afterwards. It is faster to create indexes on an empty table.

SAS Scalable Performance Data (SPD) Server Extensions to Base SAS

You can access SAS Scalable Performance Data (SPD) Server by using an SQL pass-through CONNECT statement or you can issue a SAS LIBNAME statement. After connecting to SAS Scalable Performance Data (SPD) Server, you can run SAS DATA steps, SAS procedures, or PROC SQL statements.

The documentation in the SAS Scalable Performance Data (SPD) Server Administrator's Guide and the SAS Scalable Performance Data (SPD) Server User's Guide furnish syntax and examples that use SAS Scalable Performance Data (SPD) Server extensions to Base SAS language. Most of your existing SAS programs will work in SAS Scalable Performance Data (SPD) Server with only minor modifications.

SAS Scalable Performance Data (SPD) Server extensions to the Base SAS language include:

- new LIBNAME statement options
- SAS Scalable Performance Data (SPD) Server SQL pass-through syntax
- new table options
- new macro variables
- parallel WHERE-clause processing
- parallel group-by processing
- BY-data grouping
- parallel index creation
- PROC SPDO, an operator interface procedure.

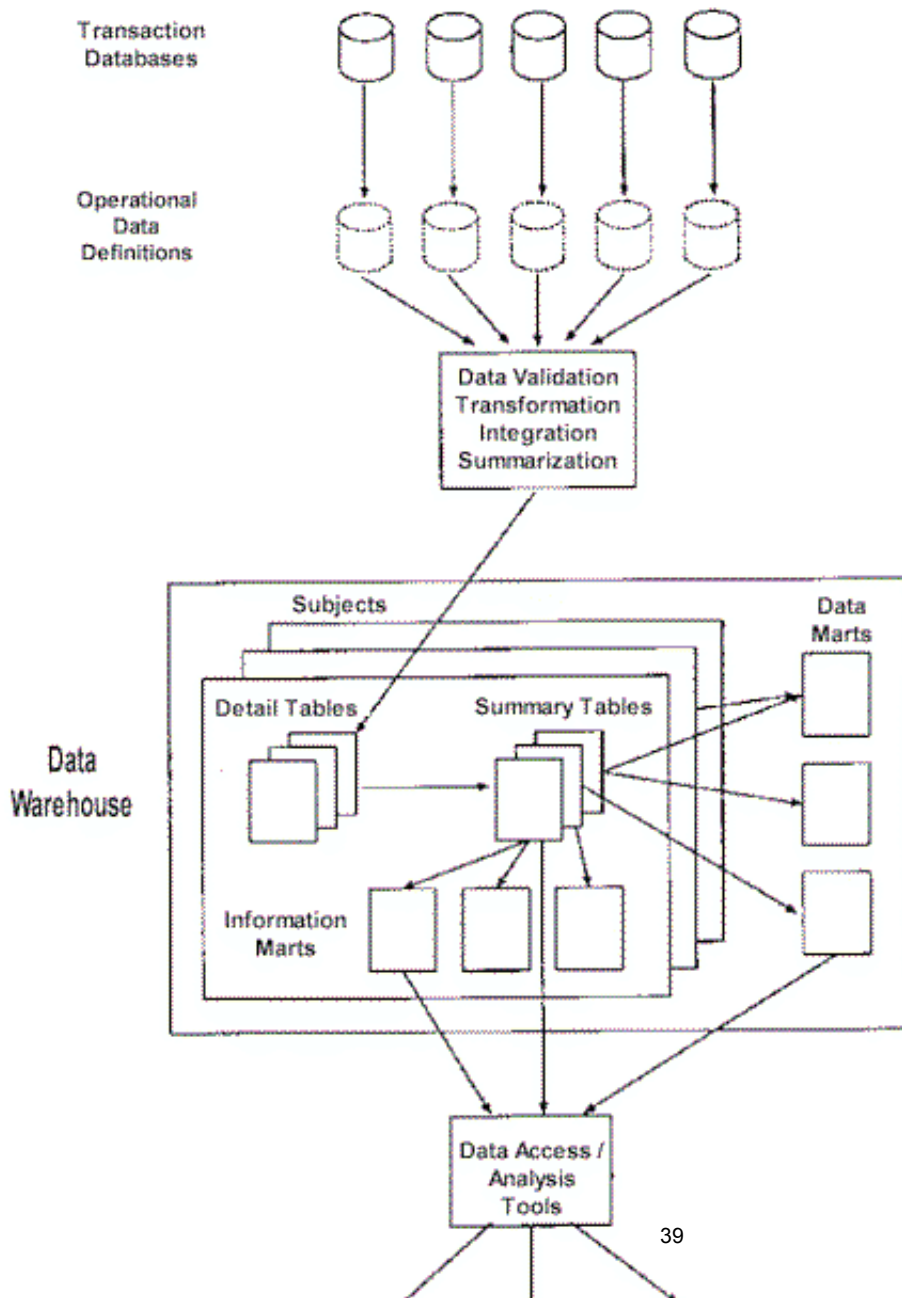
Using SAS Scalable Performance Data (SPD) Server with Data Warehousing


SAS Scalable Performance Data (SPD) Server offers SAS Data Warehousing customers an excellent facility to store data. Using component files and partitioning, SAS Scalable Performance Data (SPD) Server alleviates large table constraints such as device or directory size limits. SAS Scalable Performance Data (SPD) Server can perform storage services on a reliable and relatively inexpensive machine.

Besides providing efficient, economical storage, SAS Scalable Performance Data (SPD) Server can deliver the enhanced processing capabilities users need to manage and query data in a warehouse. SMP processing furnishes the machine's horsepower to parallel-process huge tables. SAS Scalable Performance Data (SPD) Server also offers multiple access, domain protection, and table locking: these features enable Data Warehouse users to secure and access their shared SAS Scalable Performance Data (SPD) Server.

Figure 1.5

Data Warehouse With Large Data Stores





Within a Data Warehouse, there are several *data stores* (repositories for data). Three stores are of interest above: Detail Tables, Summary Tables, and Data Marts. Organizations often store transactions that are up to 90 days old in a Detail store, transactions that are up to a year old in a Summary store, and additional data 'snapshots' in Data Marts. The three data stores share a common requirement -- they must maintain hundreds of gigabytes of data.

To perform queries, Data Warehouse users can use the SAS System with SAS syntax or PROC SQL syntax. Alternatively, the software supports use of other vendors' applications that allow pass-through SQL and comply with other non-SAS connection standards. In brief, SAS Scalable Performance Data (SPD) Server can contribute significantly to objectives for a Data Warehouse: to deliver low-cost, relevant, machine-independent, and timely information to users throughout the organization.

Copyright © 2008 by SAS Institute Inc., Cary, NC, USA. All rights reserved.

Connecting to SAS Scalable Performance Data (SPD) Server

- [Introduction](#)
 - [SAS and SPD Server Tables](#)
 - [SAS Data Libraries](#)
 - [Temporary LIBNAME Domains](#)
 - [SPD Server Resource Security](#)
 - [UNIX File Security](#)
 - [ACL File Security](#)
 - [Accessing SPD Server from a SAS Client](#)
 - [SQL Pass-Through Facility](#)
 - [LIBNAME Access](#)
 - [LIBNAME Options](#)
 - [Connect to a Specified SPD Server Host](#)
 - [SPDSHOST= Macro Variable](#)
 - [Validate the Client User ID](#)
 - [Manage Server Network Traffic](#)
 - [Additional LIBNAME Options](#)
 - [LIBNAME Example Statements](#)
- [SPD Server Table Options](#)
 - [Options to Enhance Performance](#)
 - [Options for Other Functions](#)
- [SPD Server Macro Variables](#)
 - [Macro Variables and Corresponding Table Options](#)
 - [Summary of SPD Server Macro Variables](#)
 - [Variable for a Client and Server Running on the Same UNIX Machine](#)
 - [Variable for Compatibility with the Base SAS Engine](#)
 - [Variables for Miscellaneous Functions](#)
 - [Variables for Sorts](#)
 - [Variables for WHERE Clause Evaluations](#)
 - [Variables That Affect Disk Space](#)
 - [Variables to Enhance Performance](#)

[Introduction](#)

All SAS users should read the Help section on [Accessing and Creating SPD Server Tables](#) to review the

methods that they can use to access SPD Server. These methods include LIBNAME statements and SQL pass-through statements. Syntax statements and options are provided for each method, as well as useful table options and macro variables.

SAS and SPD Server Tables

SPD Server tables have different physical structures than SAS tables. In a general discussion, a SAS table can also refer to an SPD Server table. If the context is specific, for example, an SPD Server command, then the reference is specific. A SAS table refers to the Base SAS format; an SPD Server table refers to the SPD Server format.

Using SPD Server and SAS together, you can

- convert tables from the Base SAS format to the SPD Server format
- convert tables from the SPD Server format to the Base SAS format
- create a new SPD Server table
- read, query, append to, update, sort, and index SPD Server tables.

SAS Data Libraries

The term 'SAS data library' refers either to a collection of SAS files or SPD Server files. For SPD Server, a *SAS data library* is a collection of one or more directories that specify the location of stored SPD Server files. A data library has a primary file system. This is the directory an SPD Server administrator defines for the LIBNAME domain when it is set up. Optionally, a data library can have other directories for separation of SPD Server component files.

An SPD Server data library can contain the following LIBNAME domain files:

- SPD Server tables
- SPD Server indexes
- SPD Server catalogs
- SPD Server ACL files
- SPD Server utility files, such as a VIEW, an MDDB, etc.

Temporary LIBNAME Domains

SPD Server allows you to create temporary LIBNAME domains that exist only for the duration of the LIBNAME assignment. Using this capability, SPD Server users can create space analogous to the SAS WORK library. To create a temporary LIBNAME domain, use the SPD Server LIBNAME statement option, TEMP=YES.

When you end your SPD Server session, all the data objects, including tables, catalogs, and utility files in the

TEMP=YES temporary domain are automatically deleted. This is similar to how the SAS WORK library functions.

SPD Server Resource Security

SPD Server provides two levels of data security: UNIX file security and ACL file security. ACL file security enforces SPD Server permissions with SPD Server user IDs and Access Control Lists (ACLs).

UNIX File Security

The software enables ACL file security by default. While ACL file security is strongly recommended, the default can be changed. Only an SPD Server administrator can change the default file security setting. When a SPD Server administrator specifies the NOACL option, all clients for SPD Server obtain the SPD Server user ID 'anonymous'. There is no SPD Server security in effect. SPD Server tables are then secured only by the UNIX file protections that are currently in force.

When UNIX file security controls SPD Server file access, it validates on the user ID associated with SPD Server. Which UNIX user ID is associated with SPD Server? The UNIX ID associated with SPD Server is the UNIX ID of the user that brings up the server. Suppose an SPD Server administrator brings up the SPD Server host machine, using his SPD Server administrator's account named SPDSADMIN. When any SAS client connects to this SPD Server host, they will only be able to read files that have UNIX read permissions set for the SPDSADMIN user. As a result, SAS clients that are connected to this SPD Server host must write all files in a directory created by SPDSADMIN that also has write permission set for SPDSADMIN. SPDSADMIN will own all files written in this directory.

How is security maintained? The SPD Server administrator can set up the SPD Server LIBNAME domain directories such that only the administrator has appropriate read and write access to those directories.

It is possible for a site to give different UNIX permissions to a group of users. To do this, an SPD Server administrator must bring up another SPD Server using a different UNIX user account. (Bringing up a different SPD Server affects only the new SPD Server files created, not existing SPD Server files.)

ACL File Security

UNIX file security alone is not adequate for many installations. For more complex workplace environments, SPD Server provides a finer level of controls, called ACL file security. ACL file security is used by default for SPD Server LIBNAME domains. SPD Server always enforces ACL file security unless an SPD Server administrator specifies the NOACL option when bringing up a Server.

To understand ACL file security, you must know how SPD Server user IDs work. The SPD Server administrator assigns each approved SPD Server user an ID, a password, a level of data authorization, and, optionally, membership in up to five ACLGROUPS. (The SPD Server user ID 'anonymous' does not require a

password.)

Once your SPD Server UserID has been created, you and the SPD Server administrator can use PROC SPDO to create ACLs that grant or deny other users access to an SPD Server table. The documentation chapter on [Accessing and Creating SAS Scalable Performance Data Server Tables](#) explains how to use the PROC SPDO operator interface to secure SPD Server resources.

Accessing SPD Server from a SAS Client

SQL Pass-Through Facility

SPD Server SQL pass-through processing supports an associated proxy process for each new client (via the name server). The proxy issues SQL pass-through requests. To connect to an SPD Server SQL server from a SAS session, you must submit a CONNECT statement that specifies the SASSPDS engine and SPD Server options, and then issues the SQL commands.

For example:

```
PROC SQL;
connect to sasspds
  (dbq='mydomain'
   host='namesvrID'
   serv='5555'
   user='neraksr'
   passwd='siuya');
select *
  from connection
  to sasspds
  (select * from employee_info);
disconnect from sasspds;
quit;
```

LIBNAME Access

A logical name, or *libref*, is a name for the data library that you associate with an SPD Server domain during a SAS job or session. Once a libref is assigned, SPD Server allows you to read, create, or update files in the data library if you have the appropriate access to the data library.

A libref is valid only for the current SAS job or session. Librefs can be referenced repeatedly during a valid job or session. SAS does not limit the number of librefs that you can assign during a session. Once you define a libref, it is most commonly used as the first element in two-level SAS file names: *LibraryName.TableName*.

The library name, or libref, identifies where the SPD Server can find or store the file.

The documentation chapter on [Accessing and Creating SAS Scalable Performance Data Server Tables](#) contains several [SQL pass-through examples that use librefs](#). The following example is a libref used with LIBNAME access to an SPD Server.

Example: A LIBREF Used with LIBNAME Access

The statement below creates the table TRAVEL and stores it in a permanent SAS data library with the libref ANNUAL.

```
data annual.travel;
```

Below is a LIBNAME statement that associates a libref, the SASSPDS engine, and an SPD Server domain.

```
libname mydatalib sasspds 'mydomain'  
  host='namesvrID'  
  serv='5555'  
  user='neraksr'  
  passwd='siuya';
```

```
LIBNAME libref SASSPDS <'SAS-data-library'> <SPD Server-options>;
```

Use the following arguments:

libref

a name that is up to eight characters long and that conforms to the rules for SAS names.

SASSPDS

the name of the SPD Server engine.

'SAS-data-library'

the logical LIBNAME domain name for an SPD Server data library on the host machine. The name server resolves the domain name into the physical path for the library.

SPD Server-options

one or more SPD Server options.

LIBNAME Options

You must supply the SASSPDS engine name to access SPD Server LIBNAME domains with a LIBNAME statement. You must also specify one or more SPD Server options. The syntax for an SPD Server option is

```
<SPD Server-option>=<value>;
```

SPD Server-option

a keyword to name the option.

value

a value expected by the keyword.

Option values in a LIBNAME statement enable the engine to initiate, manage, and tailor a client session. This section *summarizes* LIBNAME options and groups them by function.

Connect to a Specified SPD Server Host

To connect to a host, SPD Server needs the network node name for the SPD Server host machine or the IP address of the server machine, and the port number of a name server. SPD Server provides the following options to locate a name server using a named service.

SERVER=

specifies a node name for an SPD Server host machine and a port number for the name server running on the machine.

HOST=

specifies a node for an SPD Server host machine and a port number for the name server running on the machine.

Both options have the same function. SERVER= arguments are compatible with SAS/SHARE software. HOST= arguments support FTP conventions. The HOST option allows a node to be an IP address (for example, 123.456.76.1); the SERVER option requires a network node name.

SPDSHOST= Macro Variable

If you create a SAS macro variable named SPDSHOST= or an environment variable named SPDSHOST=, whenever a LIBNAME statement does not specify an SPD Server host machine, SPD Server will look for the value of SPDSHOST= to identify the server.

```
%let spdshost=samson;  
libname myref sasspds 'mylib'  
    user='yourid'  
    password='swami';
```

The first statement assigns the SPD Server host SAMSON to the macro variable SPDSHOST. Therefore, a subsequent LIBNAME statement does not need to name the host server again.

Validate the Client User ID

SPD Server uses the name server to secure its domains. SPD Server uses ACL file security to secure domain resources. If ACL file security is enabled, the software grants access in the following hierarchy:

- using the permissions that belong to the UNIX ID that is associated with the SPD Server
- using the permissions that belong to the SPD Server user ID.

You can use SQL pass-through and LIBNAME options to specify the identify of an SPD Server user. SPD Server uses a special ID table to validate user IDs and passwords. The following LIBNAME options identify a client:

ACLGRP=

specifies one of up to five ACL groups that the user may belong to.

ACLSPECIAL=

grants special privileges to an SPD Server user who is previously set up as *special* (ACLSPECIAL=YES is defined for the user in the password file.) Special privileges override other ACL restrictions that apply to resources in the domain.

CHNGPASS=

prompts a client user to change his or her SPD Server password.

NEWPASSWORD= or NEWPASSWD=

specifies a new password for an SPD Server client user.

PASSWORD= or PASSWD=

specifies a password to validate an SPD Server client user.

PROMPT=

prompts for a password to validate an SPD Server client user.

PASSTHRU=

specifies implicit SQL pass-through options for an SPD Server client user.

USER=

specifies the SPD Server user ID.

Table 2. 1

User ID Options When ACL File Security Is

Enabled

User=	Password= or Prompt=	Grants Access To ...
Required unless the SAS client process has a User ID, that is, not a Windows client. Submitted values for User= are validated against the SPD Server User ID Table.	Required and validated against the SPD Server User ID Table.	Resources that you create within the SPD Server LIBNAME domain and in other resources that are not excluded by ACLs or by UNIX file permissions.

[Table 2.2](#)

User ID Options When UNIX File Security Only Is

Enabled

User=	Password= or Prompt=	Grants Access To ...
Not required. The SPD Server User ID under UNIX file security only is "anonymous".	Not required with "anonymous" User ID.	All resources within the LIBNAME domain granted by UNIX permissions for the SPD Server's UNIX ID.

[Manage Server Network Traffic](#)

If your SPD Server installation uses the same physical machine to run your SPD Server client process and your SPD Server host services, you can use the two following SPD Server options to improve client / server network traffic:

NETCOMP=

compresses the data stream in an SPD Server network packet.

UNIXDOMAIN=

uses UNIX domain sockets for data transfer between the client and the SPD Server.

Additional LIBNAME Options

BYSORT=

performs an implicit sort when a BY clause is encountered.

DISCONNECT=

specifies when to close network connections between the SAS client and the SPD Server. This may be after all librefs are cleared or at the end of a SAS session.

ENDOBS=

specifies the end row (observation) in a user-defined range.

NOSASSORT=

ignores an explicit PROC SORT statement.

STARTOBS=

specifies the start row (observation) in a user-defined range.

TRUNCWARN=

Suppresses hard failure on NLS transcoding overflow and character mapping errors. When using the TRUNCWARN=YES LIBNAME option, data integrity may be compromised because significant characters can be lost in this configuration. The default setting is NO, which causes hard read/write stops when transcode overflow or mapping errors are encountered. When TRUNCWARN=YES, and an overflow or character mapping error occurs, a warning is posted to the SAS log at data set close time if overflow occurs, but the data overflow is lost.

LIBNAME Example Statements

Example 1

Example 1 creates the libref MINE, associates it with the SASSPDS engine, and specifies the SPD Server LIBNAME domain GOLDMINE. Values for the SPD Server options specify to

- locate the server machine FASTCPUS and use the default service SPDSNAME to get the port number of the name server
- validate the SPD Server user EXPLORER
- prompt for EXPLORER's old SPD Server password
- change the password.

```
libname mine sasspds 'goldmine'  
  user='explorer'  
  host='fastcpus'  
  prompt=yes  
  chngpass=yes;
```

Example 2

Example 2 represents the first LIBNAME statement that was made for the SPDSDATA domain. It creates the libref MYLIB, associates MYLIB with the SASSPDS engine, and specifies the SPD Server libname domain SPDSDATA. Values for the SPD Server options specify to

- locate the server machine HEFTY and use the named service SPDSNAME to get the port number of the name server.
- validate the SPD Server user ID **camills** and account password of **escort**.
- store data file partitions in the directories MAINDATA on device DISK1, MOREDATA on device DISK2, and MOREDATA on device DISK3. This example implies that the metadata and index partitions for tables are stored in the primary file system, that is, the path set up by the SPD Server administrator for SPDSDATA.

```
libname mylib sasspds 'spdsdata'  
  server=hefty.spdsname  
  user='camills' password='escort'  
  datapath=( '/disk1/maindata'  
             '/disk2/moredata'  
             '/disk3/moredata' );
```

[SPD Server Table Options](#)

SPD Server table options specify processing actions that apply only to a specific table. When you use a LIBNAME statement, you should specify the options in parentheses next to the table name. If you use an SQL pass-through statement, use brackets to specify the options next to the table name.

Options to Enhance Performance

BYNOEQUALS=

specifies the index output order of table rows with identical values for the BY column.

NETPACKSIZE=

controls the size of an SPD Server network data packet.

SEGSIZE=

sizes the segment for index files associated with an SPD Server table.

Options for Other Functions

BYSORT=

performs an implicit sort of a given table when a BY clause is encountered and there is no index available.

ENDOBS=

specifies the end row (observation) number in a user-defined range.

STARTOBS=

specifies the start row (observation) number in a user-defined range.

SORTSIZE=

specifies the amount of memory (in number of bytes, not Kbytes or Mbytes) that SPD Server is able to allocate in order to complete a sorting request. The SORTSIZE= table option declared must be less than the global sortsize parameter specified in the spdsserv.parm server parameter file.

VERBOSE=

details all indexes associated with an SPD Server table. This option also provides other information, such as who is the table owner and the ACL group.

SPD Server Macro Variables

You can use global macro variables in SPD Server to simplify your work. Global macro variables use default values set by the SPD Server software and operate in the background. You can make global changes to the values of macro variables in your code by specifying a new the default setting for the specified variable. The new default setting is applied to all macro variables in the code that you submit to SPD Server. You can also override the setting for a single macro variable by using a table option to change the setting for only the specified table.

The default macro variable values automate sophisticated processing decisions. The default settings furnish good performance. However, top performance often requires intelligent changes to some macro variable default settings. When you make changes to the macro variable default settings, you should attempt to find the best processing opportunity for the type of data that you have.

Learning the best way to set SPD Server macro variables and options takes time. Sometimes, performance testing is the only way to determine if changing a setting improves processing performance. Performance testing is time well spent. After you quantify performance parameters under various macro variable settings, you can customize SPD Server so that it solves your real business or data problems with maximum efficiency.

Each SPD Server installation is different. You may want to change many values, or just a few default values. When you make changes, you will find macro variables are friendly, flexible and easily to manipulate.

Use a %LET statement to change macro variable values. You can place the macro variable assignment anywhere in the open code of a SAS program except data lines. The most convenient place to put your %LET statements to initialize macro variables is in your autoexec.sas file or at the beginning of a program. The macro variable assignment is valid for the duration of your session or the executing program. Macro variable values remain in effect until they are changed by a subsequent assignment.

Assignments for macro variables with YES|NO arguments must be entered in uppercase (capitalized).

Because the SPD Server macro variables operate behind the scenes, you cannot query SPD Server to find out the status of a macro variable. SAS does not 'know' about the status of macro variables. If you want to see *which* SPD Server macro variables are in effect, or their default values, you can use PROC SPDO.

[Macro Variables and Corresponding Table Options](#)

When you need to apply the action to a single table that a macro variable applies globally to all tables, you should use a table option instead of the macro variable setting. A table option is more selective because you can turn the macro variable function on or off for a single table.

[Summary of SPD Server Macro Variables](#)

This section *summarizes* the SPD Server macro variables and groups them by the function of their default value.

Variable for a Client and Server Running on the Same UNIX Machine

SPDSCOMP=

specifies to compress the data when sending a data packet through the network.

Variable for Compatibility with the Base SAS Engine

SPDSBNEQ=

specifies the output order of table rows with identical values in the BY column.

Variables for Miscellaneous Functions

SPDSEOBS=

specifies, when processing a table, the end row (observation) number in a user-defined range.

SPDSSOBS=

specifies, when processing a table, the start row (observation) number in a user-defined range.

SPDSUSAV=

specifies, when appending to tables with unique indexes, to save rows with non-unique (rejected) keys to a separate SAS table.

SPDSUSDS=

returns the name of a hidden SAS table generated by the SPD Server which stores rows with identical (non-unique) table values.

SPDSVERB=

specifies when executing a PROC CONTENTS statement to provide more details that are specific to SPD Server indexes that are associated with the table. Examples of information include ACL information, index information, PARTSIZE= value, and others.

SPDSFSAV=

specifies to retain the table if an abnormal condition is encountered during a table-creation operation. (Normally SAS closes and deletes these tables.)

SPDSEINT=

specifies disconnect behavior for the SQL pass-through EXECUTE() statement.

Variables for Sorts

SPDSBSRT=

specifies for the SPD Server to perform a sort whenever it encounters a BY clause, and there is no index available.

SPDSNBIX=

specifies whether to turn BY-sorts with an index on or off.

SPDSSTAG=

specifies whether to use non-tagged or tagged sorting for PROC SORT or BY processing.

Variables for WHERE Clause Evaluations

SPDSTCNT=

specifies the number of threads to be used for WHERE clause evaluations.

SPDSEV1T=

specifies whether the data returned from WHERE clause evaluations that utilize an index should be in strict row (observation) order.

SPDSEV2T=

specifies whether the data returned from WHERE clause evaluations that do **not** utilize an index should be in strict row (observation) order.

SPDSWDEB=

specifies when evaluating a WHERE expression, whether WHINIT, the WHERE clause planner, should display a summary of the execution plan.

SPDSIRAT=

controls, when WHERE clause processing with enhanced bitmap indexes, whether to perform segment candidate pre-evaluation.

Variables That Affect Disk Space

SPDSCMPF=

specifies to add a number of bytes to a compressed block as growth space.

SPDSDCMP=

specifies to compress SPD Server tables on the disk.

SPDSIASY=

specifies, when creating multiple indexes on an SPD Server table, whether to create the indexes in parallel.

SPDSSIZE=

specifies the size of an SPD Server table partition.

Variables to Enhance Performance

SPDSNETP=

sizes a buffer in server memory for the network data packet.

SPDSSADD=

specifies whether to apply a single row, or multiple rows at a time, when appending to a table.

SPDSSYRD=

specifies whether to perform data streaming when reading a table.

SPDSAUNQ=

specifies whether to abort an append if uniqueness is not maintained.

Copyright © 2008 by SAS Institute Inc., Cary, NC, USA. All rights reserved.

Accessing and Creating SAS Scalable Performance Data (SPD) Server Tables

- [Introduction](#)
- [Using a LIBNAME Statement to Access SPD Server](#)
 - [Example: Issuing an Initial LIBNAME Statement](#)
 - [The Client Session](#)
- [Managing Large SPD Server Files](#)
 - [Initial Setup of SPD Server LIBNAME Domain Storage](#)
 - [Effect of the Administrator Option, ROPTIONS=](#)
 - [Explicit or Default Storage Paths](#)
 - [Understanding SPD Server Component Storage](#)
 - [Forced Partitioning of the Data Component](#)
 - [Importance of the First Metadata Partition](#)
 - [Using Path Options for Large Table Storage](#)
 - [Example 1: Specify an Explicit Initial Set of Paths](#)
 - [Example 2: Specify a Subsequent LIBNAME Statement to Add Paths](#)
- [Interchanging SPD Server and SAS File Formats](#)
 - [Migrating Tables between SAS and SPD Server](#)
 - [Example 1: Create a SAS Table from an SPD Server Table](#)
 - [Example 2: Convert from SAS to SPD Server Format](#)
- [The SQL Pass-Through Facility](#)
 - [Accessing Data Using the SQL Pass-Through Facility](#)
 - [SQL Pass-Through Statements](#)
 - [CONNECT Statement](#)
 - [DISCONNECT Statement](#)
 - [EXECUTE Statement](#)
 - [CONNECTION TO Statement](#)
 - [Example 1: Using SAS PROC SQL to Connect to an SQL Server](#)
 - [Example 2: Nested SQL Pass-Through](#)
- [Creating a New Table](#)
 - [Example - Creating a New Table Using Pass-Through Statements](#)
 - [Example - Creating a New Table with a LIBNAME Statement](#)

[Introduction](#)

This documentation chapter describes how to access SPD Server using SAS and a SPD Server SQL pass-through facility or SAS LIBNAME statement. The chapter also demonstrates typical data tasks on an SPD Server host. Finally, it discusses how to secure SPD Server resources using PROC SPDO. (Power users who have special privileges should see [Using PROC SPDO](#), Special and Privileged OPER Commands.)

Note: For readability, the SPD Server SQL pass-through facility is shortened here to "SQL pass-through facility," unless the context requires a more explicit reference. Similarly, when the chapter references a name server, it is the Scalable Performance Data Server name server.

Using a LIBNAME Statement to Access SPD Server

It is not necessary to understand all possible LIBNAME and table options to initiate an SPD Server client session. There are only a few required elements which are shown in the example below. The LIBNAME statement should specify

- the local library reference (libref)
- the required engine name SASSPDS
- a valid domain name that is registered to the name server and defined to the SPD Server host
- the name server host's name
- the user ID
- password access, either through the [PROMPT=YES](#) switch or using the [PASSWD](#) keyword. (The [PROMPT=YES](#) approach is recommended for security reasons.)

Example: Issuing an Initial LIBNAME Statement

```
libname market sasspds 'mktdata' host='sunone'  
user='user id' prompt=yes;
```

This example specifies the libref "market," the engine name SASSPDS, the LIBNAME domain "mktdata," and the name server host called "sunone." It identifies an SPD Server user "user id" and is configured to prompt the user for a password. Alternately, but less recommended, is

```
libname market sasspds 'mktdata' host='sunone'  
user='user id' passwd='beemer';
```

The only difference between this and the previous example is the password specification. Here the password "beemer" is recurred into the LIBNAME statement. This method can be used for batched SPD Server jobs that run unattended.

The Client Session

Successfully issuing the LIBNAME statement or SQL pass-through statement(s) initiates an SPD Server client session. The client session operates using a combination of up to four distinct components:

SPD Server Name Server

The name server acts as a "traffic cop" and provides a central point of control between clients and SPD Server hosts. The name server maintains a list of LIBNAME domains associated with each SPD Server host. Client sessions will always connect to an SPD Server host through a name server. The name server resolves the submitted LIBNAME domain name (a logical entity) to a physical path (usually a UNIX or Windows directory). The name server then connects you to the SPD Server serving the domain without requiring you to know physical addresses. An SPD Server administrator sets up the LIBNAME domains in a parameter file for SPD Server which then registers its domains with the name server.

SPD Server Host

Each SPD Server host controls security access to the domain resources it manages. When an SPD Server host starts up, it registers its LIBNAME domains with the name server. Clients may only connect to an SPD Server host through a name server -- direct connections between clients and SPD Server hosts are not permitted. The SPD Server host validates the client user ID and password (passed in the LIBNAME statement), launches the system process (client proxy) for each client, and grants access to the appropriate SPD Server domain.

SQL Server

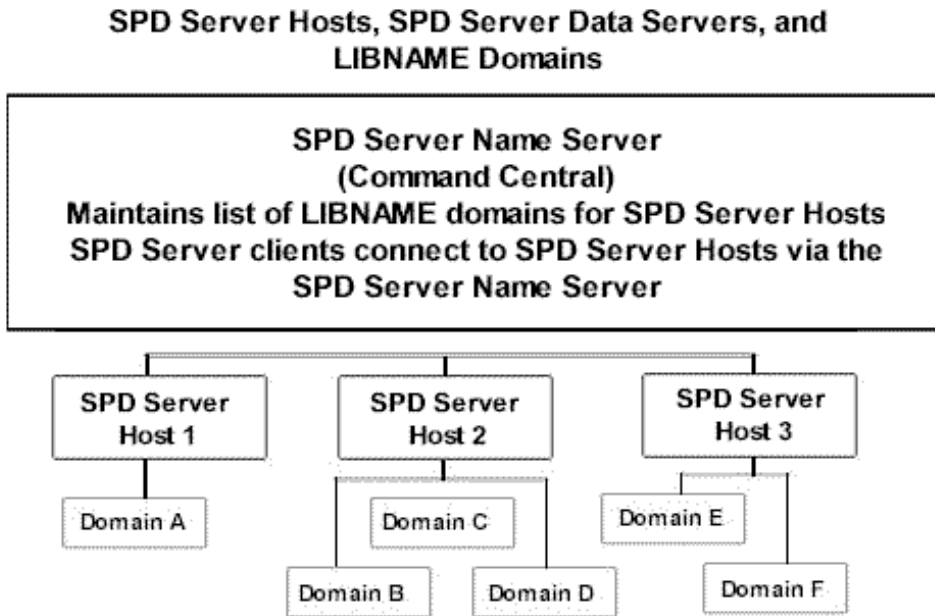
The SQL server parses and processes the pass-through SQL syntax submitted by the SAS client.

SPDSSNET Server

The SPDSSNET server enables access between clients without SAS software and SAS Scalable Performance Data Server. The SPDSSNET server runs as a stand-alone process on either the client or SPD Server host machine. It acts as a bridge between

the SAS ODBC driver and the SPD Server host. SPDSSNET also can be used with JDBC drivers and HTMSQL used with Web Servers. SPDSSNET can run multiple processes concurrently and perform parallel processing.

Figure 3.1 SPD Server Hosts, SPD Server Name Servers, and LIBNAME Domains



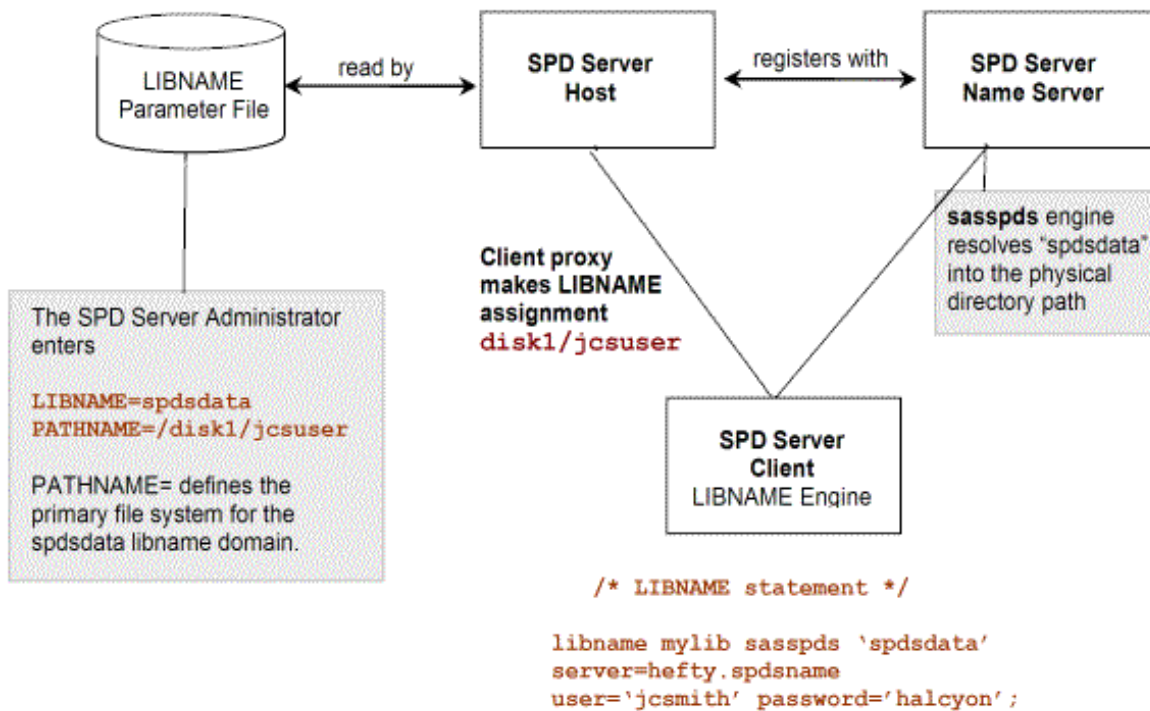
Managing Large SPD Server Files

Leaving aside performance issues, managing large files is a matter of file storage and disk space. Optimally, an SPD Server administrator will manage storage space for SPD Server LIBNAME domains. In this case, you do not need to consider storage issues -- SPD Server does the work for you. The Help section on [Optimizing SPD Server Performance](#) contains more detail on managing large SPD Server files.

Initial Setup of SPD Server LIBNAME Domain Storage

Figure 3.2 reviews how an SPD Server domain is set up. An SPD Server administrator must define the name and primary path for the domain in the LIBNAME parameter file for SPD Server. The path that the administrator defines for each domain is referred to as the primary file system for that domain. The LIBNAME parameter file is read by the SPD Server at startup. The SPD Server registers the domains with the SPD Name Server. When the user issues a LIBNAME statement, the client sends a message to the SPD Name Server that will resolve the domain name to its physical directory path and also determine the SPD Server that registered the domain.

Figure 3.2 Setup of SPD Server LIBNAME Domains



The [Scalable Performance Data 4.3 Overview](#) documentation chapter discusses LIBNAME path options that allow a user to specify additional storage devices and paths for a domain. To manage their own disk space, a user must be aware of the [DATAPATH=](#), [METAPATH=](#), and [INDEXPATH=](#) options, as well as the ROPTIONS= option that the SPD Server administrator uses.

Effect of the Administrator Option, ROPTIONS=

After defining a primary file system for a domain, an SPD Server administrator can use LIBNAME parameter file options, identical to the DATAPATH=, METAPATH=, and INDEXPATH= options in the LIBNAME statement, to set up additional paths for the domain. However, the administrator can also exercise an option to restrict a the user from defining additional paths on the LIBNAME statement with the ROPTIONS= LIBNAME parameter file option. When an SPD Server administrator uses the ROPTIONS= option, the administrator's specification takes precedence over the users. More information is available in the Help section on Configuring LIBNAME Domain Disk Space in the SPD Server Administrator's Guide.

For example, assume that a user uses the DATAPATH= option to specify a path(s) to store table data for a domain, and that the SPD Server administrator also uses the DATAPATH= option, along with ROPTIONS= for that domain entry in the LIBNAMES parameter file. The user's DATAPATH= specifications are then ignored.

The administrator's use of ROPTIONS= with path options is recommended. It relieves users of the complicated task of managing disk space and avoids the need to embed physical path information in SAS programs. Instead, SAS jobs need to refer to only the logical LIBNAME, relying on ROPTIONS= embedded by the administrator to specify all of the physical information. This approach utilizes the power of the name server, allowing it to resolve path information for an SPD Server domain.

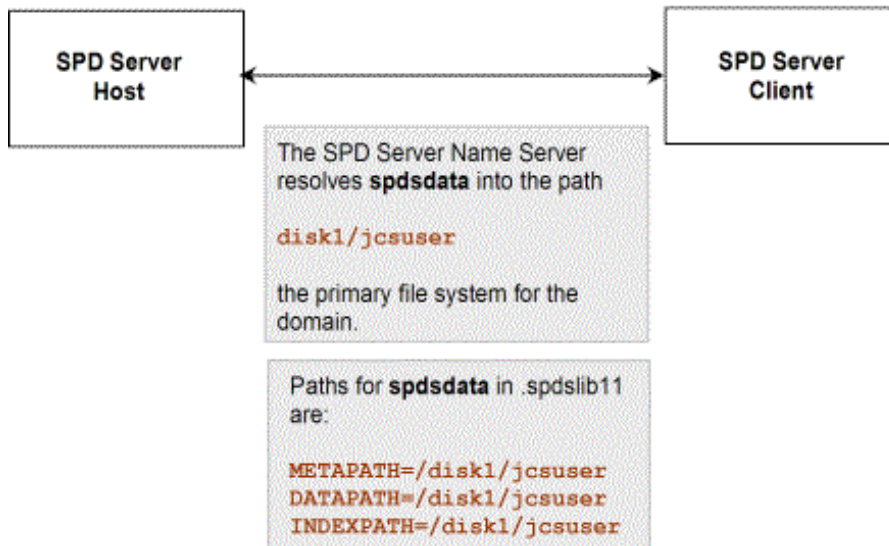
[Figure 3.3 Primary File System Default Paths](#)

```

/* First LIBNAME statement for the
domain without path options */

libname mylib sasspds 'spdsdata'
server=hefty.spdsname
user='jcsmith' password='halcyon';

```



Explicit or Default Storage Paths

You may wonder why the software offers you path options and then discourages their use? The answer is flexibility. A site may elect to allow users to manage their own disk space. While this practice is not recommended, the software allows for the possibility.

To use path options effectively, you must know that the **first LIBNAME assignment or SQL Pass-Through CONNECT statement** naming a domain establishes an initial set of paths for the domain. You can specify the paths, or the software can establish a default set. Figure 3.2 shows a default set of paths. Figure 3.4 shows an explicit initial set of paths.

The path options [METAPATH=](#), [DATAPATH=](#) and [INDEXPATH=](#) store partitions for the component files: metadata, data, and indexes. Subsequent LIBNAME assignments augment the path list created by the initial LIBNAME assignment. That is, SPD Server appends each new path assignment to any prior list for the component.

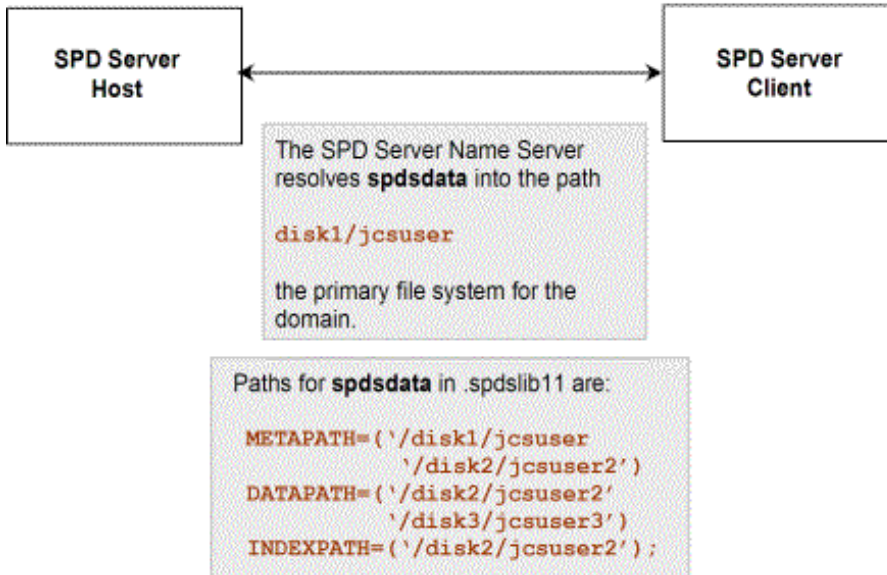
[Figure 3.4 Explicit Initial Set of Paths](#)


```

/* First LIBNAME statement for
domain with path options */

libname mylib sasspds 'spdsdata'
server=hefty.spdsname
user='jcsmith' password='halcyon'
METAPATH=('/disk2/jcsuser2')
DATAPATH=('/disk2/jcsuser2'
          '/disk3/jcsuser3')
INDEXPATH=('/disk2/jcsuser2');

```



In summary, unless you or an SPD Server administrator specify an initial set of paths, the software uses the domain's primary file system in the LIBNAMES parameter file for the default path set. As you will learn in the next section, the default path set may not be ample for large tables nor provide optimal performance.

Understanding SPD Server Component Storage

Earlier, you learned that the software creates a list of paths for storage of table files in an SPD Server domain, but file partition storage was not discussed. This section focuses on using path options when an SPD Server administrator has not used [ROPTIONS=](#).

Minimally, each table consists of a metadata component and a data component. Each component file is composed of one or more partition files on disk. The software requires that the *first metadata partition* reside in the primary file system, that is, the path defined for the domain by an SPD Server administrator. Other metadata partitions can overflow to additional paths specified using the [METAPATH=](#) option.

If no paths are specified for index and data components by the [INDEXPATH=](#) or [DATAPATH=](#) options, the software stores these partitions in the primary file system too. If other paths are specified, the software stores the initial partition for these classes in the first path with available space. (Unlike metadata partitions, data and index partitions do not have to start in the primary file system.) A partition can expand until the path fills up; remaining partitions then overflow to the next path with available space, and so on. (See [Figure 3.5](#).)

Forced Partitioning of the Data Component

To improve parallel processing of various operations involving full-table scans (for example, WHERE-clause evaluations without indexes or SQL GROUP-BY evaluations) the SPD Server allows you to force creation of data component partitions at fixed-size intervals. To specify the size interval, use the [PARTSIZE=](#) table option. By default, the SPD Server sets PARTSIZE= to 16 megabytes.

See the documentation chapter on [SAS Scalable Performance Data Server Table Options](#) for details.

The SPD Server uses the collection of file systems that you specify with the `DATAPATH=` option to distribute partitions in a cyclic (round-robin) fashion. But, instead of creating partitions until the first file system is full, the SPD Server randomly chooses a file system from the `DATAPATH=` list for the first partition, then sequentially assigns partitions to successive file systems in the `DATAPATH=` list. The software continues to cycle through the file system set, as many times as needed, until all data partitions for the table are stored. Assume that you specify

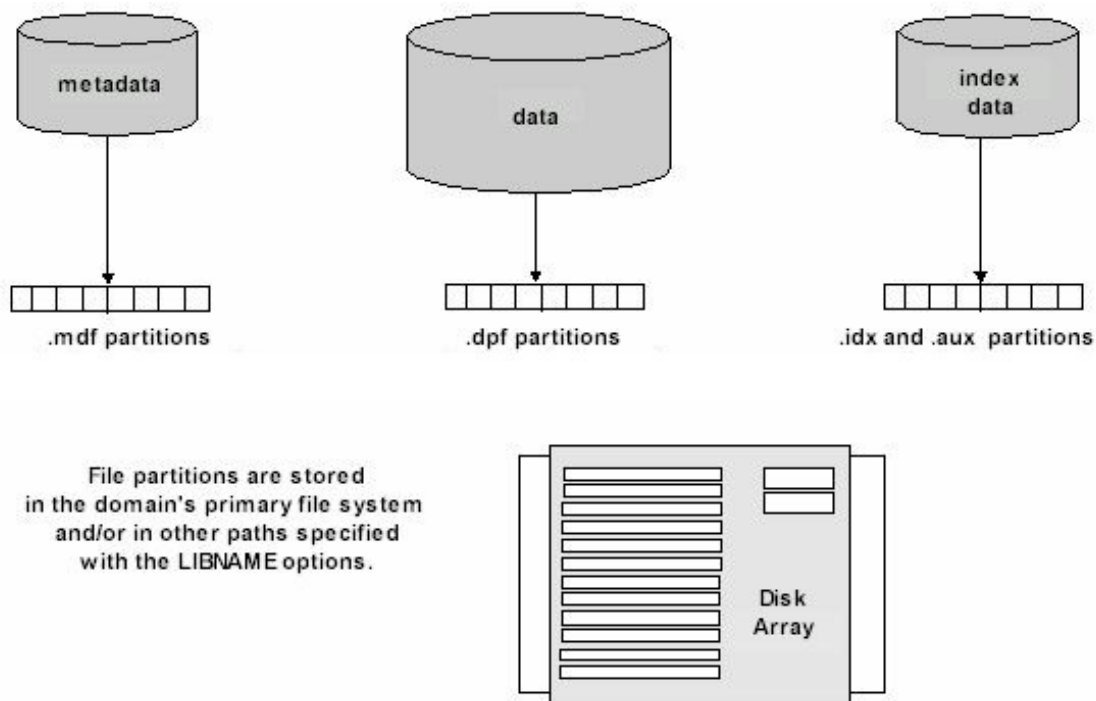
```
DATAPATH=(' /data1' '/data2')
```

Subsequently, you store your BIGONE table into the domain. SPD

uses random placement of data partitions in the `DATAPATH=` list, so the first first BIGONE partition may be stored in either the `/data1` or the `/data2` directory. Subsequent partitions will alternate between the `/data1` and `/data2` directories, and so on.

If you set `PARTSIZE=0`, SPD Server uses the `DATAPATH=` file systems strictly as overflow space. That is, it creates partitions in the first file system, up to the file size limit of your operating system. Then, when the first file system is full, it proceeds to the second file system, etc.

Figure 3.5 SPD Server Component Storage



What happens when you issue the first `LIBNAME` statement for a domain but do not specify path options? If your tables are small, most likely the primary file system is probably adequate. However, if you store large tables, the primary file system can fill up quickly. How do you know when the primary file system is full? SPD Server will return an error message when you perform an append operation on an existing table or create a new table in the domain.

Importance of the First Metadata Partition

If the primary file system is full, you may issue a subsequent `LIBNAME` statement specifying additional paths. This allows a data

append to an existing table but may not allow creation of a new table in the domain. The reason why the new paths did not solve the create failure may not be obvious. The answer is the software cannot store the first metadata file partition because the primary file system is still full. What is the create failure solution? Either free space in the primary file system or have the SPD Server administrator create a new LIBNAME domain.

Using Path Options for Large Table Storage

If you must manage your table storage, anticipate disk space for large tables. Use the LIBNAME path options with the first LIBNAME statement for the domain. Store data and index partitions using the [DATAPATH=](#) and [INDEXPATH=](#) options on a different storage device than the primary file system. This reserves the primary file system for metadata files.

Example 1: Specify An Explicit Initial Set of Paths

SITEUSR1 issues the first LIBNAME statement for the MYLIB domain. By default, the domain's primary file system is used to store metadata partitions but another device MYDISK30 and directory SITEUSER is specified to store the data and index partitions. (The SPD Server administrator created the primary file system for MYLIB.)

```
/* I anticipate the primary file system for the MYLIB domain      */
/* is ample for metadata files, but I will use MYDISK30          */
/* to store my data and index partitions.                        */
libname myref sasspds 'mylib'
    datapath=('/mydisk30/siteuser')
    indexpath=('/mydisk30/siteuser')
    server=husky.spdsname
    user='siteusr1' prompt=yes;
```

Example 2: Specify A Subsequent LIBNAME Statement to Add Paths

SITEUSR1 issues a subsequent LIBNAME statement for the MYLIB domain specifying additional paths for the data and index partitions. The user is storing very large tables so two storage devices (and directories) for data are listed, and a third device for indexes associated with the tables is listed.

```
/* I noticed today MYDISK30 is getting full.                    */
/* I am adding MYDISK31 for possible overflows.                 */
libname expand sasspds 'mylib'
    datapath=('/mydisk31/siteuser' '/mydisk32/siteuser')
    indexpath=('/mydisk33/siteuser')
    server=husky.spdsname
    user='siteusr1' prompt=yes;
```

The software appends the new paths listed to the prior list for each component type. The entire path list that **spdslib11** now maintains is

```
datapath=('/mydisk30/siteuser' '/mydisk31/siteuser' '/mydisk32/siteuser')
indexpath=('/mydisk30/siteuser' '/mydisk33/siteuser')
```

How does SPD Server use the path list? It stores partitions of the data components for MYLIB tables in the specified data paths. (How the software uses the paths depends upon the value of the [PARTSIZE=](#) option.) For index components, it stores the files in the first path listed until the space is filled, then it proceeds to fill the next path listed.

Interchanging SPD Server and SAS File Formats

Migrating Tables between SAS and SPD Server

Many organizations use SPD Server when they discover there is a need for more "horsepower" dealing with large SAS tables. As a result, there are many instances where it is handy to be able to move SAS tables into SPD Server format, and vice versa. Fortunately, SPD Server was designed with ease of table conversion in mind. The examples below illustrate the flexibility built into SPD Server and the ease of table conversion between the SAS and SPD Server systems.

Example 1: Create a SAS Table from an SPD Server Table

To create a SAS table from an SPD Server table, issue a LIBNAME statement but do not specify the engine SASSPDS. Your program will then create a Base SAS table. (Later, if you decide to use SPD Server capabilities, you can convert the SAS table to the SPD Server format. Conversion is easy: interchange table formats using the SAS System's COPY procedure. See Example 2.)

```
/* Create local racquets data set. */
libname local '/u/sasdemo/local';

data local.racquets;
  input racquet_name $20. @22 weight_oz @28 balance $2.
      @32 flex @36 gripsize
      @42 string_type $3. @47 retail_price @55 inventory_onhand;
  datalines;
Filbert VolleyMaster 10.5 HL 5 4.5 syn 129.95 5
Solo Queensize      10.9 HH 6 5.0 syn 130.00 3
Perkinson AllCourt  11.0 N 5 4.25 syn 159.99 12
Wilco Specialist     8.9 HL 3 5.0 nat 287.50 1
;
```

Example 2: Convert from SAS to SPD Server Format

SITEUSR1 makes a libref SPORT, associates SPORT with the SPD Server engine SASSPDS, and points to the CONVERSION_AREA domain on an SPD Server host server named HUSKY. User SITEUSR1 uses a default named service SPDSNAME to locate the port number of the name server and requests a prompt for the password.

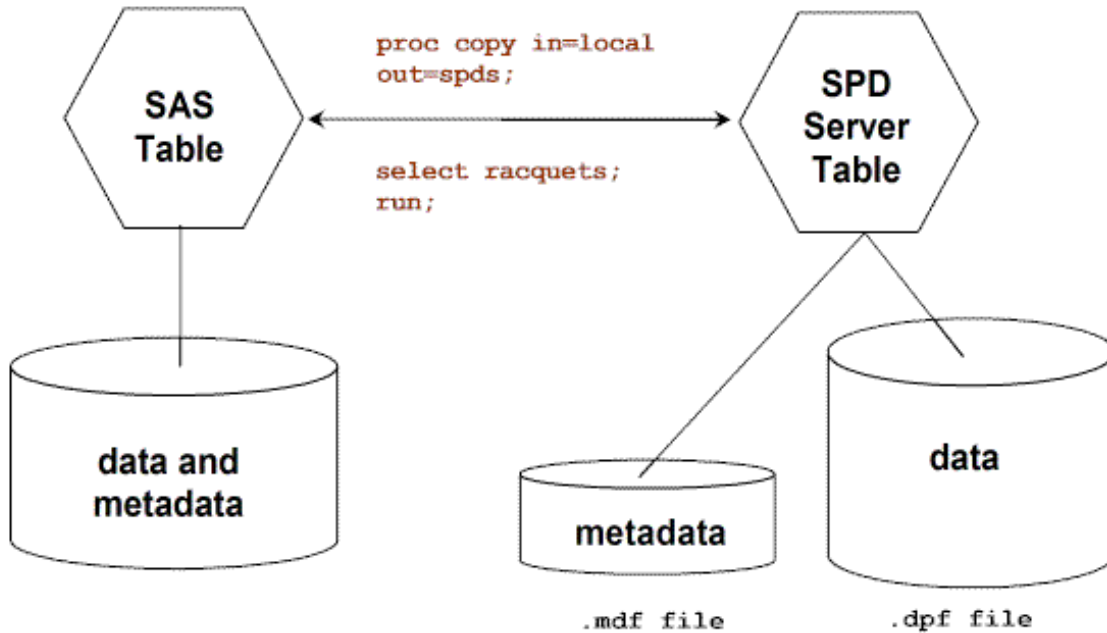
The PROC COPY statement inputs the SAS table LOCAL.RACQUETS and outputs the SPD Server table SPORT.RACQUETS to the CONVERSION_AREA domain. After the PROC COPY statement executes, the SAS table becomes two SPD Server table component files. (See Figure 3.6.)

```
/* Copy existing SAS table to the SPD Server format. */
libname sport sasspds 'conversion_area' server=husky.spdsname

      user='siteusr1' prompt=yes;

proc copy in=local out=sport;
  select racquets;
run;
```

Figure 3.6 PROC COPY Converts a SAS Table to an SPD Server Table



The SQL Pass-Through Facility

SPD Server uses pass-through SQL commands to access and manipulate data. What does this mean? Enabling pass-through SQL functionality provides SPD Server clients with a new way to establish a connection with an SPD Server host or direct load from an external database such as Oracle. Users now have broader data access in the SPD Server environment and growing connectivity to external databases using the SPD Server engine.

The [SQL Syntax Reference Guide](#) documentation chapter provides additional detailed reference information on using SPD Server SQL syntax.

Accessing Data Using the SQL Pass-Through Facility

The SQL pass-through facility is another access method allowing SPD Server to connect to an SQL server and manipulate data. An overview of the steps is presented here, and followed with examples. These are the major steps for using SQL pass-through:

1. Establish a connection from an SPD Server client using a [CONNECT](#) statement.
2. Send SPD Server SQL statements using the [EXECUTE](#) statement.
3. Retrieve data SQL query with the [CONNECTION TO](#) component in a SELECT statement's FROM clause.
4. Terminate the connection using the [DISCONNECT](#) statement.

SQL Pass-Through Statements

CONNECT Statement

Specifies the SAS I/O engine that will provide the SQL pass-through access.

Syntax

CONNECT TO *dbms-name* < **AS** *alias* >(*dbms-args*);

Use the following arguments:

dbms-name (required)

Specifies the name of the engine.

When running SAS and PROC SQL, you must specify **sasspds** to obtain SQL pass-through **to** an SPD Server SQL Server. You must specify **spdseng** to obtain SQL pass-through **from** an SPD Server SQL server. The later examples show CONNECT statements specifying these engines.

AS alias (optional)

Specifies an alias or logical name for a connection.

When specifying an alias to identify the connection, use a string without quotes. Then refer to this logical name in subsequent SQL pass-through statements.

Note: The alias must specify the connection that will execute the statement.

Example - Using an Alias

```
execute(...) by alias
```

or

```
select * from connection to alias(...)
```

dbms-args (required and/or optional arguments)

Identifies the SQL server and data source. The following *dbms-args* arguments are for the SPD Server engines, **sasspds** and **spdseng**. SPD Server SQL uses the following simple syntax: *Keyword=Value*

DBQ=libname-domain (required)

Specifies the primary SPD Server LIBNAME domain for the SQL pass-through connection.

The name that you specify is identical to the LIBNAME domain name that you used when making a SAS LIBNAME assignment to **sasspds**. Use single or double quotes around the specified value.

HOST=name-server-host (optional)

Specifies a node name or IP address for a name server that is currently running.

Use single or double quotes around the specified string. If you do not specify a name, the software uses the current value of the SAS macro variable **spdshost** to determine the node name.

SERVICE=name-server-port (optional)

SERV=name-server-port (optional)

Specifies the network address (port number) for a name server that is currently running.

Use single or double quotes around the specified value. If you do not furnish a port number for the name server, the software determines the port address from the named service **spdsname** in the **/etc/services** file.

USER=SPD Server user ID (required on Windows but not UNIX)

Specifies an SPD Server user ID to access an SPD Server SQL Server. Use single or double quotes around the specified value.

PASSWORD=password (required)

PASSWD=password (required, or use PROMPT=YES, unless USER='anonymou')

Specifies an SPD Server user ID password to access an SPD Server. (This value is case sensitive.) Normally you would not specify a password in text files that others can view. More likely you would use this argument in batch jobs that are protected by file system permissions, prohibiting others from reading the job files.

PROMPT=YES (required, or use PASSWD or PASSWORD=, unless USER='anonymou')

Specifies a password prompt to access an SPD Server SQL server. This value is case sensitive.

DISCONNECT Statement

Disconnects you from your DBMS source.

Syntax

DISCONNECT FROM [*dbms-name* | *alias*];

Description

When you are finished with a PROC SQL connection, you must disconnect from the DBMS source. This automatically occurs when you exit the PROC SQL procedure. You can, however, explicitly disconnect from the DBMS by using the DISCONNECT statement.

Use the arguments:

dbms-name

the name specified in the CONNECT statement that established the connection.

alias

the alias value specified in the CONNECT statement that established the connection.

EXECUTE Statement

The EXECUTE statement is part of the pass-through SQL facility. It allows the user to use specific SQL statements during a pass-through connection. Before using the EXECUTE statement, the user must first establish a connection using the CONNECT statement. After a user has created a pass-through connection, use EXECUTE to submit valid SQL statements (except the SELECT statement).

Syntax

EXECUTE (*SQL statement*) **BY** [*dbms-name* / *alias*];

Use the following arguments:

(*SQL statement*)

A valid SQL statement passed for execution (except SELECT statements). This argument is required and must be enclosed within parentheses.

dbms-name (required, or use *alias*)

Identifies the DBMS to which you want to direct the SQL statement. Note that *dbms-name* must be preceded by the keyword **BY**.

alias (optional, or use *dbms-name*)

Specifies an optional alias used in the CONNECT statement.

CONNECTION TO Statement

CONNECTION TO is an SQL pass-through component that can be used in a SELECT statement's FROM clause as part of the *from*-list. The CONNECTION TO component enables you to make pass-through queries for data and to use that data in a PROC SQL query or table. PROC SQL treats the results of the query like a virtual table.

Syntax

CONNECTION TO *dbms-name* (*SQL-query*)

Use the following arguments:

dbms-name (required)

If you have a single connection, *dbms-name* is the *dbms-name* specified in your CONNECT statement. If you have multiple connections, use the alias specified in the AS clause of the CONNECT statement.

(*SQL-query*)

The (*SQL-query*) specifies the SQL query you want to send. Your SQL query cannot contain a semicolon because that represents the end of a statement to SPD Server. Character literals are limited to 32,000 characters. Be sure your SQL query is enclosed in parentheses.

alias (optional)

Specifies an optional alias used in the CONNECT statement.

Example 1: Using SAS PROC SQL to Connect to an SQL Server

To connect from a SAS session to an SQL server, in this example the SPD Server's SQL Server, execute a CONNECT statement. After making the connection, the first execute statement creates a table EMPLOYEE_INFO with three columns, EMPLOYEE_NO, EMPLOYEE_NAME, and ANNUAL_SALARY. The second execute statement inserts an observation into the table where EMPLOYEE_NO equals "1" and EMPLOYEE_NAME equals "The Prez".

The subsequent FROM CONNECTION TO statement retrieves all the records from the new EMPLOYEE_INFO table. (In this example, that would be the single observation inserted by the second execute statement.) The DISCONNECT statement terminates the data source connection.

```
PROC SQL;
connect to sasspds
  (dbq='mydomain'
   host='workstation1'
   serv='spdsname'
   user='me'
   passwd='noway');
execute (create table employee_info
  (employee_no num, employee_name char(30),
   annual_salary num) by sasspds;
execute (insert into employee_info
  values (1, 'The Prez')) by sasspds;
select * from connection to sasspds
  (select * from employee_info);
disconnect from sasspds;
quit;
```


Example 2 - Nested SQL Pass-Through

SPD Server pass-through access can be nested. Nesting allows access to data stored on two different networks or network nodes.

In the example that follows, we nest SQL pass-through from the current local network host DATAGATE to access the EMPLOYEE_INFO table, which is available at the PROD host on a remote network. (Our example presumes that we have user access to PROD.)

```
proc sql;
connect to sasspds (dbq='domain1'
                  host='datagate' serv='spdsname'
                  user='usr1' passwd='usr1_pw');
execute (connect to spdseng (dbq='domain2'
                            host='prod' serv='spdsname'
                            user='usr2' passwd='usr2_pw') by sasspds;
select * from connection to sasspds(
    select * from connection to spdseng(
        select employee_no, annual_salary
        from employee_info));
execute (disconnect from spdseng) by sasspds;
disconnect from sasspds;
quit;
```

Creating a New Table

One of the SPD Server's strengths lies in the ability to create, manipulate, and query very large tables. As a rule of thumb, client users generally choose not to store massive tables locally because of their sheer size. The following code examples assume that users will create and store large tables on the SPD Server host.

Example - Creating a New Table Using Pass-Through Statements

First, connect from a SAS session to an SQL server, in this example the SPD Server's SQL Server, then execute a CONNECT statement. After making the connection, the first execute statement creates a table LOTTERYWIN with two columns, TICKETNO and WINNAME. The second execute statement inserts an observation into the table where TICKETNO equals "1" and NAME equals "Wishu Weremee."

The subsequent FROM CONNECTION TO statement retrieves all the records from the new LOTTERYWIN table. (In this example, that would be the single observation inserted by the second execute statement. The DISCONNECT statement terminates the data source connection.

```
proc sql;
connect to sasspds (dbq='mydomain'
                  host='workstation1' serv='spdsname'
                  user='me' passwd='luckyones');
execute (create table lotterywin
        (ticketno num, winname char(30))) by sasspds;
execute (insert into lotterywin
        values (1, 'Wishu Weremee')) by sasspds;
select * from connection to sasspds
        (select * from employee);
disconnect from sasspds;
```

```
quit;
```

Example - Creating a New Table with a LIBNAME Statement

SITEUSR1 creates a new SPD Server table CARDATA.OLD_AUTOS on the server.

```
libname cardata sasspds 'conversion_area' server=husky.5105
      user='siteusr1' prompt=yes;

/* Create the table CARDATA.OLD_AUTOS on the SPD Server host. */

data cardata.old_autos;
  input year $4. @6 manufacturer $12. model $12. body_style $5.
  engine_liters @39 transmission_type $1. @41 exterior_color
  $10. options $10. mileage conditon;

datalines;

1966 Ford      Mustang      conv  3.5  M  white      00000001 143000 2
1967 Chevrolet Corvair      sedan 2.2  M  burgundy  00000001  70000 3
1975 Volkswagen Beetle      2door 1.8  M  yellow    00000010  80000 4
1987 BMW       325is      2door 2.5  A  black     11000010 110000 3
1962 Nash     Metropolitan conv  1.3  M  red       00000111 125000 3
;
```

Copyright © 2008 by SAS Institute Inc., Cary, NC, USA. All rights reserved.

Indexing, Sorting, and Manipulating SAS Scalable Performance Data (SPD) Server Tables

- [Introduction](#)
- [Indexing a Table](#)
 - [The SPD Index](#)
 - [Creating SPD Indexes Examples](#)
 - [Creating SPD Indexes from a DATA Step](#)
 - [Creating SPD Indexes from PROC DATASETS](#)
 - [Creating SPD Indexes Using SQL](#)
 - [Creating SPD Indexes Using Pass-Through SQL](#)
 - [Using VERBOSE= to See Index Information](#)
- [Using PROC SORT with SPD Server](#)
 - [Example Using Implicit SPD Server BY Clause Sort](#)
 - [Example Using PROC SORT](#)

[Introduction](#)

This chapter describes and provides examples on indexing, sorting, and manipulating SPD Server tables on an SPD Server host.

[Indexing a Table](#)

SPD Server provides a single SPD index type that efficiently indexes tables of varying size and data distributions. The SPD Server SPD index optimally supports queries that require global table views (such as queries that contain BY Clause processing and SQL joins), or queries which require segmented views (such as parallel processing of WHERE-clause statements).

[The SPD Index](#)

The SPD index maintains two views of the index values, a global view and a segmented view. The global view is maintained using a unique global B-tree that has a single entry for each discrete value. The segmented view is maintained by the data for each value in the global B-tree, which includes a list of segments that contain the value, and for each segment a bitmap that identifies which rows in the segment contain the value. The global view is maintained in the SPD index **.hbx** file, and the segmented data is maintained in the SPD index **.idx** file.

For queries that require a global view, SPD Server searches the hybrid global B-tree for a particular value. The segment lists are scanned for the value, then the bitmaps from each segment containing the value are read. SPD Server uses the bitmap to locate and retrieve the observations for that segment. This type of query returns results sorted first by value and then by observation number. This sorting is optimal for BY Clause processing and SQL joins.

A parallel WHERE-clause on a table that is indexed is done in two phases. The first phase, pre-evaluation, uses the SPD indexes to build a list of segments that satisfy the query. The list drops segments from the WHERE-clause scan queue when those segments contain no data in the clause range. As more and more segments are *excluded* from the scan queue, the benefit of the pre-evaluation phase increases proportionally. The second phase in the evaluation launches threads which read an index in parallel. Each thread queries a particular segment of the index, using information from the pre-evaluation phase. Using the SPD index, the thread reads the segment bitmap. The per-segment bitmaps identify the segment rows which satisfy the query for that particular column. If you include more than one indexed column in the WHERE-clause, SPD Server retrieves the per-segment bitmaps for each column in parallel (as are the segments for each column). After retrieving all the bitmaps for each column of the segment, SPD Server determines which rows satisfy the query, and returns those segment rows to the client. The multi-threaded per-segment queries begin execution at the same time, and their finishing order varies and can not be reasonably predicted. As a result, the overall order of the results cannot be guaranteed when you are using this type of query. See the documentation chapter on [Understanding Whinit - the Data Server WHERE Clause Planner](#) for a more detailed description on using indexed columns with WHERE-clause evaluations.

When a table is modified due an append or update, all SPD indexes on the table are updated. Updating the index can potentially fragment the per-value segment lists or cause some disk space to be wasted. A highly fragmented SPD index can negatively impact the performance of queries that utilize the index. In this case, you should reorganize the index to eliminate the fragmentation and reclaim wasted disk space, using the ixutil utility program. For more detailed information on reorganizing an SPD index, refer to the topic on Password Manager Utility in the SPD 4.3 Administrator's Guide.

Creating SPD Indexes Examples

This section shows how to create SPD indexes for new and existing tables.

Creating SPD Indexes from a DATA Step

```
data foo.x(  
  index=(x y=(a b)));  
x=1;  
a="Doe";
```

```
        b=20;  
run;
```

This creates SPD Server table X, then creates a simple SPD index X on column X, and a composite SPD index Y on columns (A B).

Creating SPD Indexes from PROC DATASETS

```
PROC DATASETS lib=foo;  
  modify x;  
  index create x;  
  index create y=(a b);  
quit;
```

This creates the same simple and composite SPD indexes as Example 1, assuming that the same DATA step was executed without index creation included.

Creating SPD Indexes Using SQL

```
PROC SQL;  
create index x  
  on foo.x (x);  
create index y  
  on foo.x (a,b);  
quit;
```

This creates the same simple and composite SPD indexes as Example 1, assuming that the same DATA step was executed without index creation included.

Creating SPD Indexes Using Pass-Through SQL

```
PROC SQL;  
connect to sasspds (  
  dbq="path1"  
  server=host.port  
  user='anonymous' );  
  
execute(create index x on x (x))  
by sasspds;
```

```
execute(create index y on x (a,b))
by sasspds;
quit;
```

This creates the same simple and composite SPD indexes as Example 1, assuming that the same DATA step was executed without index creation included.

Using VERBOSE= to See Index Information

There will be times when you want to see information about indexes associated with a particular table. The table option `VERBOSE=` provides details of all indexes associated with an SPD Server table. For example, if the code from Example 2 above is followed with the expression below:

```
PROC CONTENTS
  data=sports.expraqs
  (verbose=yes);
run;
```

The following will be output:

```
      Alphabetic List of Index Info:
Bitmap Index (No Global Index):          GRIPSIZE
KeyValue (Min):                          4.250000
KeyValue (Max):                          5.000000
# of Discrete values:                    3
```

Using PROC SORT with SPD Server

If you use PROC SORT with SPD Server, your table will be sorted. However, you may want to understand a few sort details to avoid surprises. Assume, for example, that you submit a PROC SORT statement to sort a table not previously indexed or sorted on the BY column.

PROC SORT takes advantage of SPD Server sorting implicitly and asserts BY Clause ordering to the SPD Server. This performs the sort on the SPD Server machine, but there will still be significant I/O between the client node and the SPD Server machine. The sorted data still makes a round trip from the server machine to the client machine and back again. Fortunately, the SQL pass-through facility in SPD Server offers an extension to the SQL language to permit a table copy and sort operation, all on the server machine.

Knowing the implications of using PROC SORT with SPD Server, how can you avoid inefficiency? The

answer is to eliminate PROC SORT statements from your SAS jobs where possible. Instead, make SAS procedures and DATA steps that require BY Clause processing use SPD Server's implicit sorts.

Example Using Implicit SPD Server BY Clause Sort

```
/* The following DATA step performs a server sort on the */
/* table column PRICE. There is no prior index for PRICE */

data _null_;
set sport.expraqs;
  by price;
  if (string='nat') then do;
    put '*' @@;
    price = price - 30.00;
  end;
put raqname @30 price;
run;
```

Example Using PROC SORT

```
/* The following PROC SORT performs a server sort on the */
/* table column MODEL. There is no prior index for MODEL */

PROC SORT
  data=inventory.old_autos
  out=inventory.old_autos_by_model;
  by model;
run;
```

Copyright © 2008 by SAS Institute Inc., Cary, NC, USA. All rights reserved.

Using SAS Scalable Performance Data (SPD) Server with Other Clients

- [Overview](#)
- [Using Open Database Connectivity \(ODBC\) to Access SAS Scalable Performance Data \(SPD\) Server Tables](#)
- [Using JDBC \(Java\) to Access SAS Scalable Performance Data \(SPD\) Server Tables](#)
- [Using htmSQL to Access SAS Scalable Performance Data \(SPD\) Server Tables](#)
- [Using SQL C API to Access SAS Scalable Performance Data \(SPD\) Server Tables](#)

This chapter describes using SAS Scalable Performance Data (SPD) Server to connect with ODBC, JDBC, htmSQL, and SQL C API clients.

Overview

Scalable Performance Data Server provides ODBC, JDBC, htmSQL, and SQL C API access to SAS Scalable Performance Data (SPD) Server data stores from all supported platforms.

SAS Scalable Performance Data (SPD) Server can read tables exported from Base SAS software using PROC COPY, and, with the proper drivers installed on the network, allows queries on the tables from client machines that do not have SAS software.

There are four possible options:

- **ODBC:** Open Database Connectivity - This is an interface standard that provides a common interface for accessing databases. Many software packages running in a Windows environment are compliant with this standard and can access data created by other software. This is a good choice if you have client machines running Windows applications, such as Microsoft Excel or Microsoft Access.
- **JDBC:** Java Database Connectivity - This option allows users with browsers to log on to a Web page and make a query. The results of the request are formatted and returned to a Web page. This makes information available across a wide range of client platforms because all you need, after installing the JDBC driver on SAS Scalable Performance Data (SPD) Server, is a Web page with some Java code, and a client machine with a Java-enabled browser.
- **htmSQL:** HyperText Markup Structured Query Language - This option allows users with browsers to log on to a Web page and make a query. The results of the request are formatted and returned to a Web page. This makes information available across a wide range of client platforms. Why? After installing the htmSQL driver in SAS Scalable Performance Data (SPD) Server, all you need is an htmSQL Web page and a client machine with a browser.
- **SQL C API:** This option allows access to SAS Scalable Performance Data (SPD) Server tables from SQL statements generated by C/C++ language applications. This access is provided in the form of a C-language run-time access library. This library provides a set of functions that you can use to write custom applications to process SAS Scalable Performance Data (SPD) Server tables and to generate new ones. This library is designed to support multi-threaded applications and is available on all supported SAS Scalable Performance Data (SPD) Server platforms.

Note: GUI interfaces may not display all return codes or error messages that the server generates.

Using Open Database Connectivity (ODBC) to Access SAS Scalable Performance Data (SPD) Server Tables

Read this section if you do not have Base SAS software on the network client, but you want to access SAS Scalable Performance Data (SPD) Server tables on the network, using an ODBC compliant program, such as Microsoft Word, Query, Excel, or Access, and you have SAS Scalable Performance Data (SPD) Server tables available for use, somewhere on the network, or Scalable Performance Data Servers and SPD SNET servers running, or client machines in a Windows environment.

- [Why Use ODBC?](#)
- [Installing ODBC Drivers on the Server](#)
- [Configuring ODBC on the Client](#)
- [Preparing your Client for ODBC Installation](#)
- [Two Types of ODBC Connections](#)
- [Primary and Secondary LIBNAME Domains](#)
- [Configuring an ODBC Data Source to Connect Directly to a SAS Scalable Performance Data \(SPD\) Server](#)
- [Configuring an ODBC Data Source for SPD SNET](#)
- [Creating a Query Using an ODBC-Compliant Program](#)

Why Use ODBC?

You have SAS Scalable Performance Data (SPD) Server tables available on your network, and one or more of the following may be true:

- You do not have Base SAS software running on the Windows client, but you need to view or change SAS Scalable Performance Data (SPD) Server tables.
- You need to view or change the SAS Scalable Performance Data (SPD) Server tables using a Microsoft spreadsheet, database or word processor.

- You need to view or change SAS Scalable Performance Data (SPD) Server tables in ways that cannot be predetermined or programmed into a Web page.
- You need to view or change SAS Scalable Performance Data (SPD) Server tables using Windows tools you are familiar with.

Installing ODBC Drivers on the Server

- Instructions for installing the ODBC driver are included in the download package.

Configuring ODBC on the Client

1. Configure an ODBC data source.
2. Make your query using a Windows program.

Figure 6.1: Configure ODBC to Connect SAS Scalable Performance Data (SPD) Server Client to SAS Scalable Performance Data (SPD) Server Host

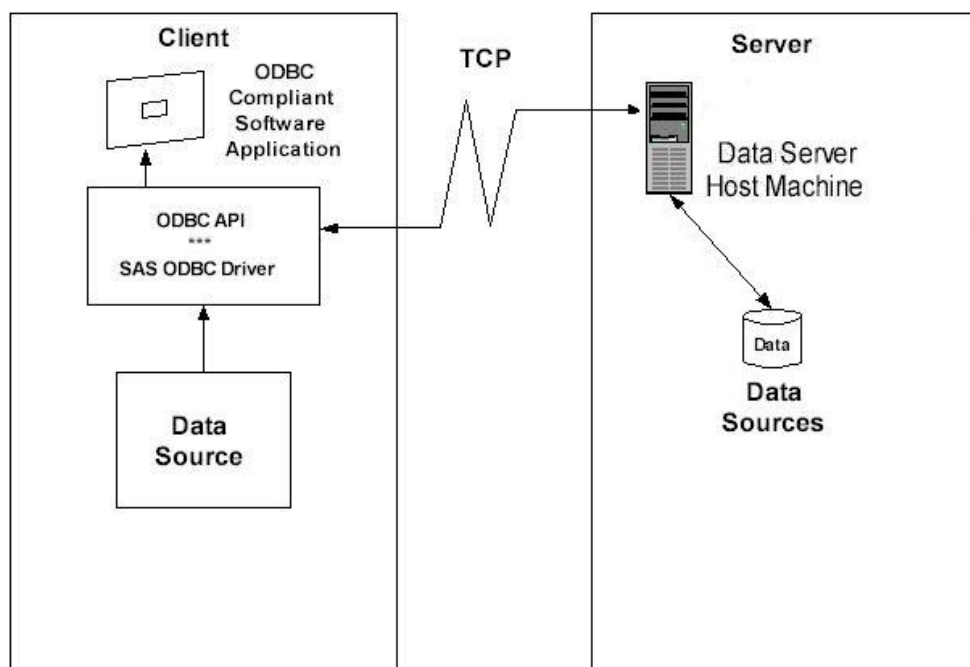
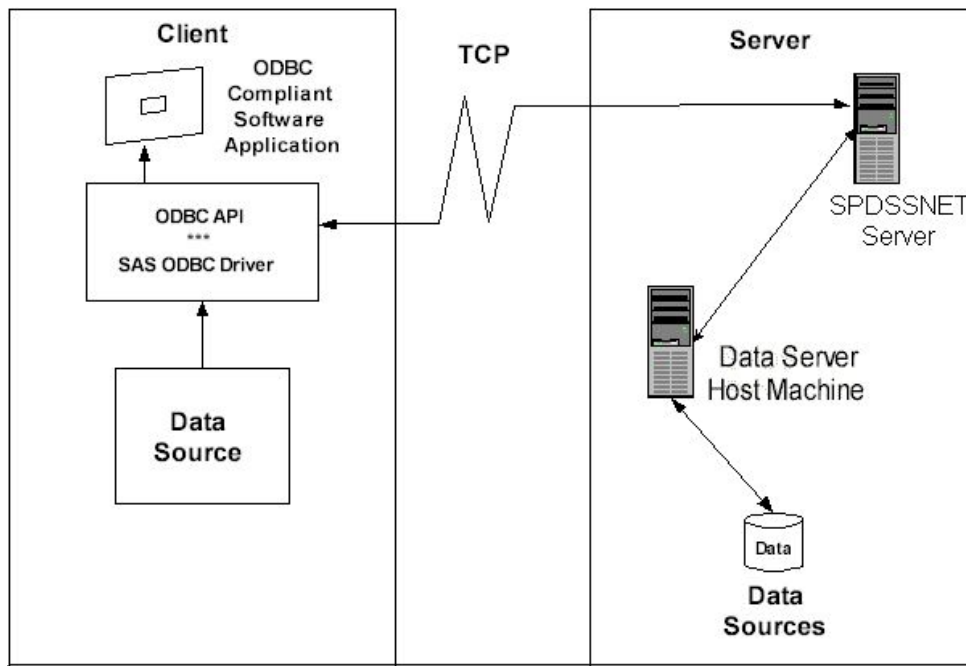


Figure 6.2: Configure ODBC to Connect SAS Scalable Performance Data (SPD) Server Client to SPD SNET Server



[Preparing your Client Machine for ODBC Installation](#)

Before you create ODBC data sources driver, you'll need the following information from your network administrator:

- o a Username and Password that is defined by a SAS Scalable Performance Data (SPD) Server administrator
- o the primary LIBNAME domain of the SAS Scalable Performance Data (SPD) Server (also called the DBQ)
- o the port number of the SPD name server (also called the SERV)
- o the machine name or IP address of the SAS Scalable Performance Data (SPD) Server Name Server (also called the HOST)
- o any secondary LIBNAME domains you wish to assign to the ODBC connection.

[Two Types of ODBC Connections](#)

With SAS Scalable Performance Data (SPD) Server software you can connect directly to a SAS Scalable Performance Data (SPD) Server without going through the SPD SNET server. Although connecting directly is the preferred method, connections via the SPD SNET server are still supported.

Note that connections via the SPD SNET server are not supported in the SAS 9 ODBC Driver software. If you intend to connect via the SPD SNET Server you must install the SAS 8 ODBC Driver.

[Primary and Secondary LIBNAME Domains](#)

When a connection to the SAS Scalable Performance Data (SPD) Server is established a primary LIBNAME domain is assigned. The primary LIBNAME domain is specified by the "DBQ" connection options parameter. Immediately after the connection is made the SAS ODBC Driver assigns the secondary LIBNAME domains which are configured through the Libraries tab of the SAS ODBC Driver Configuration window.

ODBC Connections via the SPD SNET server must have an **odbc.parm** file configured on the SPD SNET Server machine.

[Configuring an ODBC Data Source to Connect Directly to a SAS Scalable Performance Data \(SPD\) Server](#)

Once the SAS ODBC driver is installed, you will need to configure your ODBC data source. When you open the ODBC manager, you'll get a display screen that allows you to enter information that points the ODBC driver to the data on the SAS Scalable Performance Data (SPD) Server.

1. From the Windows Start button, select

Start ➔ Settings ➔ Control Panel

2. Locate the ODBC Data Sources icon and open the Microsoft ODBC Data Source Administrator . The exact location of this program depends on your version of Windows.
 3. Select the Add button, then select the SAS ODBC driver.
 4. Enter a data source name (and description if desired.)
 5. Select the Servers panel and type in your two-part server name.
 6. Click on the Configure box. The TCP Options window appears:
 - **Server Address:** Enter the network address of the machine on which the SAS Scalable Performance Data (SPD) Server is running.
 - **Server User Name:** Enter the user name as configured for a DBQ (SAS Scalable Performance Data (SPD) Server primary LIBNAME domain) on the SAS Scalable Performance Data (SPD) Server to which you will connect.
 - **Server User Password:** Enter the user password as configured for a DBQ (SAS Scalable Performance Data (SPD) Server primary LIBNAME domain) on the SAS Scalable Performance Data (SPD) Server host to which you will connect.
 - **Connection Options:** Enter the Connection Options as follows:
 - **DBQ='SPD Server primary LIBNAME domain'**, this is the SAS Scalable Performance Data (SPD) Server LIBNAME domain
 - **HOST='nameserver node name'**, this is the location of the host computer
 - **SERV='nameserver port number'**, this is the port number of the SAS Scalable Performance Data (SPD) Server name server running on the HOST.
 - Any other SAS Scalable Performance Data (SPD) Server LIBNAME options. For more information, see the User's Guide section on [LIBNAME Options](#).
 9. Click OK, then click Add, and select the Libraries panel.
 10. Enter the DBQ name of a secondary LIBNAME domain in both the Name and Host File text fields.
 11. Enter “spdseng” in the Engine text field.
 12. Follow the syntax rules for the SQL Pass-Through [LIBREF](#) statement for entering a value in the Options text field.
-

[Configuring an ODBC Data Source for SPD SNET](#)

Once the SAS ODBC driver is installed, you will need to configure your ODBC data source. When you open the ODBC manager, you'll get a display screen that allows you to enter information that points the ODBC driver to the data on the SAS Scalable Performance Data (SPD) Server.

1. From the Windows Start button, select

Start ➔ Settings ➔ Control Panel

2. Click on the ODBC icon and select the Add button.
3. Select the SAS ODBC driver.
4. Enter a data source name (and description if desired).
5. Select the Servers panel and type in the two-part server name. The second part of the server name should match the entry in the services file. In the example that follows that shows you how to edit the services file, the server name is **spdsnet**.
6. Click on the Configure box. The TCP Options window appears with four input fields that you fill:
 - **Server Address:** Enter the network address of the machine on which the SPD SNET server is running.
 - **Server User Name:** Enter the user name as configured for a DBQ (SAS Scalable Performance Data (SPD) Server primary LIBNAME domain) on the SAS Scalable Performance Data (SPD) Server to which you will connect.
 - **Server User Password:** Enter the user password as configured for a DBQ (SAS Scalable Performance Data (SPD) Server primary LIBNAME domain) on the SAS Scalable Performance Data (SPD) Server host to which you will connect.
 - **Connection Options:** Enter the connection options as follows:

- **DBQ='SPD Server primary LIBNAME domain'**: this is the SAS Scalable Performance Data (SPD) Server LIBNAME domain.
- **HOST='nameserver node name'**: this is the location of the host computer.
- **SERV='nameserver port number'**: this is the port number of the SAS Scalable Performance Data (SPD) Server name server running on the HOST.

8. Click OK, and then click Add.

[Editing the Services File on Your Machine - ODBC Details](#)

Editing the Services file is only required for ODBC connections via the SPD SNET Server.

1. Find the Services file on your Windows machine. In Windows, the Services file is usually located in

c:\windows\services

2. Open the Services file using a text editor.

3. The services file contains four columns. The rows of information may be sorted in port number order. Find the closest port number to the SAS Scalable Performance Data (SPD) Server port number, which you obtained from the network administrator (see "[Preparing for Installation](#)"). This is where you insert the new information.

4. Add an entry to the Services file, on its own line, in proper numeric order, using the following syntax:

column1	column2	column3	column4
<service name>	<port number & protocol>	<aliases>	<comment>
spdssnet	nnnn/tcp		
spdssnet=name	nnnn=port number	not required	not required
assigned to server	protocol is always /tcp		

Table 6.1: *How to Add Service Name and Port Number to the Services File*

Remember: The service name, **spdssnet** must match the server name that you used in step 6 of [Configuring an ODBC Data Source for SPD SNET](#). The port number must match the port number on which the SPD SNET server is running.

[Creating a Query Using an ODBC-Compliant Program](#)

The following instructions create a query using Microsoft Access.

1. Start the SPD SNET server.
2. Start Microsoft Access.
3. From the Microsoft Access main menu, select
File ➔ Get External Table.
4. Select **Link Table.**
5. Select **Files of Type.**
6. Select **ODBC Databases.**
7. Select the data source.

[Using JDBC \(Java\) to Access SAS Scalable Performance Data \(SPD\) Server Tables](#)

Read this information if you do not have Base SAS software on the network client, but you want to use the power of the Java programming language to query SAS Scalable Performance Data (SPD) Server tables from any client on the network that has a browser. You must have SAS Scalable Performance Data (SPD) Server tables on the

network and SAS Scalable Performance Data (SPD) Server and SPD SNET servers running on the same server as the Web server in order to use JDBC to access SAS Scalable Performance Data (SPD) Server tables.

- [Why Would I Want to Use JDBC?](#)
- [How Is JDBC Set Up on the Server?](#)
- [How Is JDBC Set Up on the Client?](#)
- [How Do I Use JDBC to Make a Query?](#)
- [JDBC Code Examples and Tips](#)
- [Limitations of Using JDBC with SAS Scalable Performance Data \(SPD\) Server](#)

Why Would I Want to Use JDBC?

You might want to use JDBC if you have SAS Scalable Performance Data (SPD) Server tables available on your network and one or more of the following is true:

- You do not have Base SAS software on the network client to process the data sets.
 - You want to distribute the information across your corporate intranet through a Web page.
 - The clients on your network are varied: UNIX boxes, Windows PCs, and workstations. One thing they might have in common is browser access to your intranet.
 - The audience for the information understands Web browsing and wants point-and-click access to the information.
 - You want to distribute the information over the World Wide Web.
 - Your planned application requires the power of the Java programming language.
-

How Is JDBC Set Up on the Server?

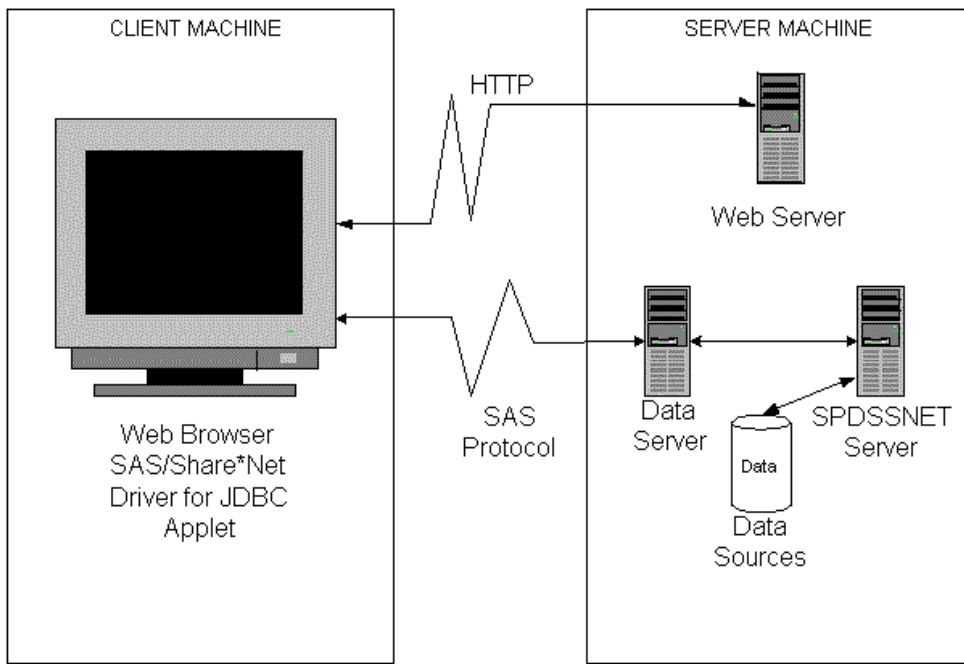
- JDBC is usually set up on the server at the time the SAS Scalable Performance Data (SPD) Server is installed. The process is covered in the SAS Scalable Performance Data (SPD) Server installation manual.
-

How Is JDBC Set Up on the Client?

The client needs a browser set up to accept Java applets, such as

- Netscape Navigator, Release 3.0 or later
- Microsoft Internet Explorer, Release 3.02 or later.

Figure 6.3: JDBC Set Up on a SAS Scalable Performance Data (SPD) Server Client



[How Do I Use JDBC to Make a Query?](#)

1. Log on to the World Wide Web and enter the URL for the Web page that contains the JDBC code.
2. Click on the desired information.
3. JDBC handles the request, formats the information, and returns the result to the Web page.

[JDBC Code Examples and Tips](#)

The following lines must be a part of the HTML file for JDBC:

```
<applet code="CLASSPATH.*.class" codebase=".." width=600 height=425>
<param name=url value="jdbc:sharenet://spdssnet_node:PORT">
<param name="dbms_options" value="DBQ='libname' HOST='host_node' SERV='NNNN'">
<param name="spdsuser" value="userid">
<param name="sharePassword" value="thepassword">
<param name="shareRelease" value="V9">
<param name="dbms" value="spds">
</applet>
```

Line 1:

- **CLASSPATH** points to the class path set up where the JDBC driver is installed.
- ***.class** is the name of the Java class that consumes all of the <PARAM name=...> lines.

Line 2:

- **spdssnet_node** is the node name of the machine on which the SPD SNET server is running.
- **PORT**=port number of the machine on which the SPD SNET server is running.

Line 3:

- **value=DBQ='libname'** is the LIBNAME domain of the SAS Scalable Performance Data (SPD) Server.
- **HOST='host_node'** is the location of the SPD SNET server.
- **SERV='NNNN'** is the port number of the name server.

Line 4:

- "spdsuser" value="userid" is the user ID that queries the SAS Scalable Performance Data (SPD) Server table.
-

Line 5:

- "sharePassword" value="thepassword" is the password of the user ID that will make the query.
-

Line 6:

- "shareRelease" value="V9" is the version of the driver you are using. This must not be altered.
-

Line 7:

- Sets the foreign database property on the JDBC driver. This means that the server is not SAS and JDBC should not create a DataBaseMetaData object. See the examples below for how to get around this.
-

Limitations of Using JDBC with SAS Scalable Performance Data (SPD) Server

- [JDBC Used with SAS Versus JDBC Used with SAS Scalable Performance Data \(SPD\) Server](#)
- [Example JDBC Query for Getting a List of Tables](#)
- [Example JDBC Query for Getting Metadata about a Specific Table](#)

JDBC Used with SAS Versus JDBC Used with SAS Scalable Performance Data (SPD) Server

SAS Scalable Performance Data (SPD) Server is treated as a foreign database. SAS Scalable Performance Data (SPD) Server clients can't query the JDBC metadata class for available tables and other metadata. Users must write their own queries to do this.

Example JDBC Query for Getting a List of Tables

(JDBC Used with SAS Scalable Performance Data (SPD) Server)

```
SELECT '' AS qual,
LIBNAME AS owner,
MEMNAME AS name,
MEMTYPE AS type,
MEMNAME AS remarks FROM dictionary.tables AS tbl
WHERE ( memtype = 'DATA' OR memtype = 'VIEW' OR memtype = 'SYSTEM TABLE' OR
        memtype = 'ALIAS' OR memtype = 'SYNONYM')
AND (tbl.libname NE 'MAPS' AND tbl.libname NE 'SASUSER' AND tbl.libname NE 'SASHELP')
ORDER BY type, qual, owner, name
```

Example JDBC Query for Getting Metadata about a Specific Table

(Your data file)

```
SELECT '' AS qual,
LIBNAME AS owner,
MEMNAME AS tname, name,
length AS datatype,
type || ' ',
length AS prec,length,
length AS scale, length AS radix, length AS nullable,label,
FORMAT FROM dictionary.columns AS tbl
WHERE memname = 'your data file'
AND (tbl.libname NE 'MAPS'
      AND tbl.libname NE 'SASUSER'
      AND tbl.libname NE 'SASHELP')
```

[Using htmSQL to Access SAS Scalable Performance Data \(SPD\) Server Tables](#)

Read this section if you do not have Base SAS software on the network client, but you want to use the point-and-click convenience of a Web page to query SAS Scalable Performance Data (SPD) Server tables from any browser-enabled client on the network. You must have SAS Scalable Performance Data (SPD) Server tables available for use, htmSQL loaded and configured on a UNIX or Windows operating system, and Scalable Performance Data Servers and SPD SNET servers running.

- [Why Would I Want to Use htmSQL?](#)
- [How Is htmSQL Set Up on the Server?](#)
- [How Is htmSQL Set Up on The Client?](#)
- [How Do I Use htmSQL to Make a Query?](#)
- [Examples of Setting Up an htmSQL Web Page](#)

[Why Would I Want to Use htmSQL?](#)

You may want to use htmSQL if you have SAS Scalable Performance Data (SPD) Server tables available on your network and one or more of the following is true:

- You do not have Base SAS software on the network client to process the data sets.
- You want to distribute the information across your corporate intranet through a Web page.
- The clients on your network are varied: UNIX boxes, Windows PCs, and workstations. One thing they might have in common is browser access to your intranet.
- The audience for the information understands Web browsing and wants point-and-click access to the data.
- You would like to use the JDBC option to extract the information but cannot permit Java applets to run on your network browsers.
- You want to distribute the information over the World Wide Web.
- Your developers are familiar with SQL and HTML.

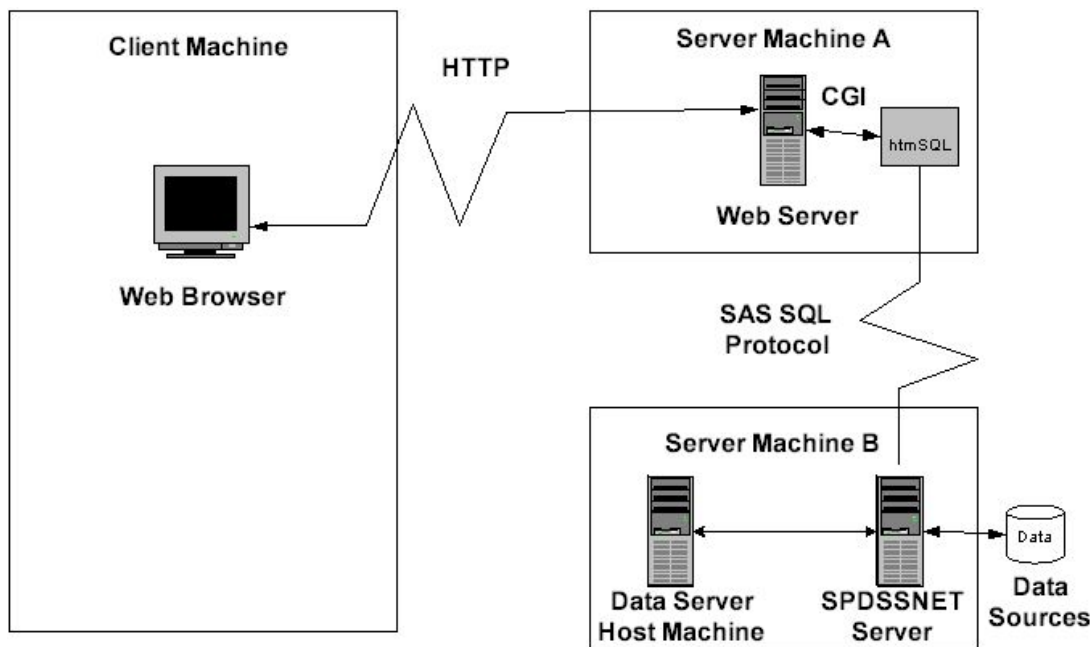
[How Is htmSQL Set Up on the Server?](#)

- htmSQL is usually set up on the server at the time the SAS Scalable Performance Data (SPD) Server is installed. The process is covered in the SAS Scalable Performance Data (SPD) Server installation manual.
- htmSQL must be installed on the Web server and you need the name of a data source that points to the SPD SNET server and to the specific LIBNAME domain that contains the SAS Scalable Performance Data (SPD) Server data you are interested in.

[How Is htmSQL Set Up on the Client?](#)

HtmSQL requires nothing more than a browser on the network or Web client.

[*Figure 6.4: htmSQL Configured on a SAS Scalable Performance Data \(SPD\) Server Client*](#)



[How Do I Use htmSQL to Make a Query?](#)

1. Log on to the World Wide Web and enter the URL for the Web page that contains the htmSQL code.
2. Click on the desired information.
3. htmSQL handles the request, formats the information, and returns the result to the Web page.

[Examples of Setting Up an htmSQL Web Page](#)

SAS Institute maintains a Web site that explains the technical details of setting up htmSQL Web pages. In some cases, there are references to the SAS/SHARE product. The rules for setting up htmSQL for either the SAS Scalable Performance Data (SPD) Server or SAS/SHARE are virtually the same.

The SAS Institute Web page for htmSQL is

<http://support.sas.com/rnd/web/intrnet/htmSQL/index.html>

[Using SQL C API to Access SAS Scalable Performance Data \(SPD\) Server Tables](#)

Read this section if you do not have Base SAS software on the network client but you want to provide your network client machines with the capability of accessing SAS Scalable Performance Data (SPD) Server tables, using SQL query methods. You must have SAS Scalable Performance Data (SPD) Server tables available for use, SAS Scalable Performance Data (SPD) Servers and SPD SNET servers running, and Network client machines capable of running C/C++ programs.

[Why Would I Want to Use SQL C API?](#)

You have SAS Scalable Performance Data (SPD) Server tables available on your network and one or more of the following may be true:

- You do not have Base SAS software on the network client to process the data sets.
- You wish to distribute the information across your corporate intranet.
- The clients on your network are varied: UNIX boxes, Windows PCs, workstations. One thing they might have in common is the ability to run C/C++ programs.
- Your developers are familiar with SQL and C/C++.

[SAS Scalable Performance Data \(SPD\) Server SQL Access Library API Reference](#) contains additional information on SQL C API.

SAS Scalable Performance Data (SPD) Server Dynamic Cluster Tables

Contents

- [Introduction to Dynamic Cluster Tables](#)
 - [Dynamic Cluster Table Structure](#)
 - [Benefits of Dynamic Cluster Tables](#)
 - [Parallel Loading](#)
 - [Fast and Economical Refreshes](#)
 - [Creating and Controlling Dynamic Cluster Tables](#)
 - [Create a Dynamic Cluster Table](#)
 - [Dynamic Cluster Table Access Control](#)
 - [Add Tables to a Dynamic Cluster](#)
 - [Undo Dynamic Cluster Tables](#)
 - [Refresh Dynamic Cluster Tables](#)
 - [Modify Dynamic Cluster Tables](#)
 - [Dynamic Cluster BY Clause Optimization](#)
 - [Dynamic Cluster BY Clause Optimization Example](#)
 - [Member Table Requirements for Creating Dynamic Cluster Tables](#)
 - [Table Attributes](#)
 - [Variable Attributes](#)
 - [Index Attributes](#)
 - [Querying and Reading Member Tables in a Dynamic Cluster](#)
 - [Unsupported Features in Dynamic Cluster Tables](#)
 - [Dynamic Cluster Table Examples](#)
 - [Create a Dynamic Cluster Table Example](#)
 - [Add Tables to a Dynamic Cluster Example](#)
 - [Undo Dynamic Cluster Table Example](#)
 - [Refresh Dynamic Cluster Table Example](#)
-

[Introduction to Dynamic Cluster Tables](#)


SPD Server is designed to meet the storage and performance demands that are associated with processing large amounts of data using SAS. As the size of the data grows, the demand to process that data increases, and storage architecture must change to keep up with business needs.

SPD Server offers dynamic cluster tables. Earlier releases of SPD Server provided a type of cluster table called the time-based partitioning table. To optimize the benefits of the clustering, the SPD Server administrator can use dynamic clusters to partition SPD Server data tables for speed and enhanced I/O processing. Clustering is performed using metadata that when combined with SPD Server functionality, provides parallel processing capabilities for loading and querying data tables. Parallel processing can accelerate performance and increase the manageability, flexibility, and scalability of very large data stores.


[Dynamic Cluster Table Structure](#)

The SPD Server dynamic cluster table can be considered as part of a hierarchy of tables with increasing sophistication:

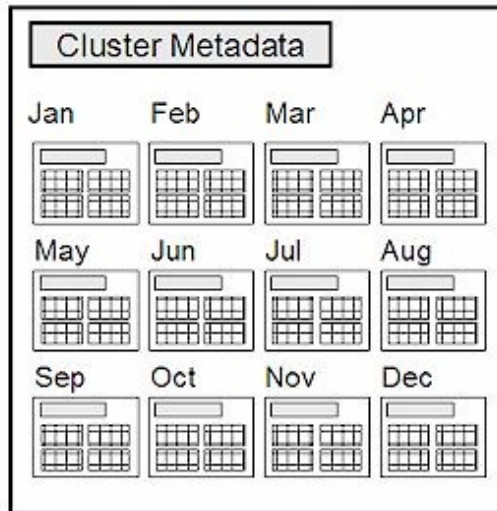
Traditional SAS Table



SPD Server Table



SPD Server Cluster Table



Traditional SAS tables are single files that contain the data descriptors and the table data. Data values are the columns, and the descriptors are the metadata that describe the column and data formatting that the table uses. If a traditional SAS table contains one or more indexes, they are stored in a separate file.

SPD Server tables use component files to store tables. One component file stores the stream of data values. Another component file stores the column and data descriptors. If you create an index for a column or a composite of columns, SPD Server creates two separate component files (a *.hbx file and a *.idx file) for each index.

SPD Server Cluster tables are virtual table structures. SPD Server cluster tables are composed of members. Each member is an SPD Server table. All members must share the same metadata formats and organization. SPD Server cluster tables use the metadata to manage the data that is contained in the members.

The SPD Server dynamic cluster table structure provides architecture that enables flexible loading and rapid storage and processing for very large data tables. Using dynamic cluster tables, loading data, removing data, and refreshing tables in very large data marts become easier and more timely. Dynamic cluster tables provide organizational features and performance benefits that traditional SAS tables and SPD Server tables do not have.

Benefits of Dynamic Cluster Tables

Organizing SPD Server data into dynamic cluster tables creates an architecture that supports parallelism, enhanced data flexibility and manageability, and significantly improved speed in robust data warehousing environments that use large and very large data tables.

For example, you can add new data or remove historical data from very large tables by accessing only the member tables that are affected by the change. You can access the individual member tables in parallel. This strategy reduces the time that is needed for the job to complete and uses simple commands. Furthermore, a complete refresh of a dynamic cluster table can occur using a fraction of the disk space that is needed to refresh a large traditional SAS or SPD Server table that contains the same amount of data.

Parallel Loading

Because dynamic cluster tables are virtual tables that consist of numerous small SPD Server tables, the architecture enables parallel loading and processing. Cluster table loads and refreshes are broken down into multiple tasks that can be performed concurrently. Separate SAS MP CONNECT jobs manage the parallel loading and processing.

The scalability of parallel loading with dynamic cluster tables depends on the scalability of the server I/O and the number of processors on the server.

Parallel loading requires multiple concurrent writes to disk. If the I/O hardware does not scale appropriately, the loading process can degrade performance.

SPD Server can create multiple indexes on the same table in parallel, and index creation is a CPU-intensive process. When sufficient processing power is available, parallel index creation in SPD Server is highly scalable. The creation process for each index is multi-threaded. A single index creation can use multiple CPUs on a server if they are available, which greatly improves performance.

Fast and Economical Refreshes

Refreshing a dynamic cluster table requires only a fraction of the disk space that a traditional SPD Server table with the same amount of data would require. The dynamic cluster table architecture allows users to refresh many large tables concurrently, while conserving disk and I/O resources. With very large traditional SAS or SPD Server tables, available disk space often limits the number of tables that can be concurrently refreshed.

In the life cycle of data warehouses, tables can be refreshed to recapture disk space when rows have been updated or deleted, or to reorder data for optimized performance. However, refreshing a table can temporarily use twice the disk space of the table itself. With very large tables, disk space can be a limiting factor when updating a data warehouse or data mart. When disk space is limited on a server, the amount of data that can be refreshed at any given time is constrained. The window of time that is required to load and refresh can become huge.

Because dynamic cluster tables can be quickly unbound into smaller SPD Server tables, refreshing dynamic cluster tables does not use twice the disk space of the original table. Instead, only twice the disk space of the largest member table in the dynamic cluster table is required.

After the dynamic cluster table is unbound, disk space equal to the first member table is required to perform a refresh. A backup of the refresh is created, and then the old version is deleted, creating more available disk space. The refresh process repeats for each successive member table until all members in the dynamic cluster table have been refreshed and updated. Then, the member tables are merged into a dynamic cluster table once again.

When a server has enough disk space and I/O resources to refresh more than one member table at a time, the benefits of parallel processing can be realized.

Creating and Controlling Dynamic Cluster Tables

Creating dynamic cluster tables in SPD Server is simple and straightforward. The following operations are associated with creating and controlling dynamic cluster tables:


- [Create a Dynamic Cluster Table](#)
- [Dynamic Cluster Table Access Control](#)
- [Add Tables to a Dynamic Cluster](#)
- [Undo Dynamic Cluster Tables](#)
- [Refresh Dynamic Cluster Tables](#)

Create a Dynamic Cluster Table

To create dynamic cluster tables in SPD Server, you must have a set of related SPD Server tables that you want to cluster, such as tables that contain monthly sales histories. The SPD Server tables that you want to cluster must all be in the same domain, and must use identical table structures (columns and indexes) and compression. However, member table partition sizes and member table owners can vary. These requirements ensure the metadata compatibility that is necessary to create dynamic cluster tables in SPD Server.

Once the related SPD Server tables are organized, a simple PROC SPDO command is used to bind the tables into a dynamic cluster table.

The following graphic represents a dynamic cluster table with 24 members. Each member table is an SPD Server table that contains monthly sales transactions:



The following code shows the PROC SPDO command syntax that is used to create dynamic cluster tables from the member tables:

```
PROC SPDO library=domain-name ;
  cluster create Sales_History
    mem=sales200301
    mem=sales200302
    mem=sales200303
    mem=sales200304
    mem=sales200305
    mem=sales200306
    mem=sales200307
    mem=sales200308
    mem=sales200309
    mem=sales200310
    mem=sales200311
    mem=sales200312
    mem=sales200401
    mem=sales200402
    mem=sales200403
    mem=sales200404
    mem=sales200405
    mem=sales200406
    mem=sales200407
    mem=sales200408
    mem=sales200409
    mem=sales200410
    mem=sales200411
    mem=sales200412
  maxslot=36 ;
quit ;
```

PROC SPDO uses a LIBRARY statement to identify the domain that contains the tables to be clustered. The **cluster create** syntax specifies the name of the dynamic cluster table to be created (Sales_History).

The **mem=** syntax identifies the members of the cluster table. The tables in the previous example represent monthly sales transactions. This example uses 24 monthly sales tables for the years 2003 and 2004.

The **maxslot=** specification specifies the maximum number of members that are allowed in the dynamic cluster table Sales_History.

The "[Dynamic Cluster Table Examples](#)" section contains more extensive code examples of creating dynamic cluster tables.

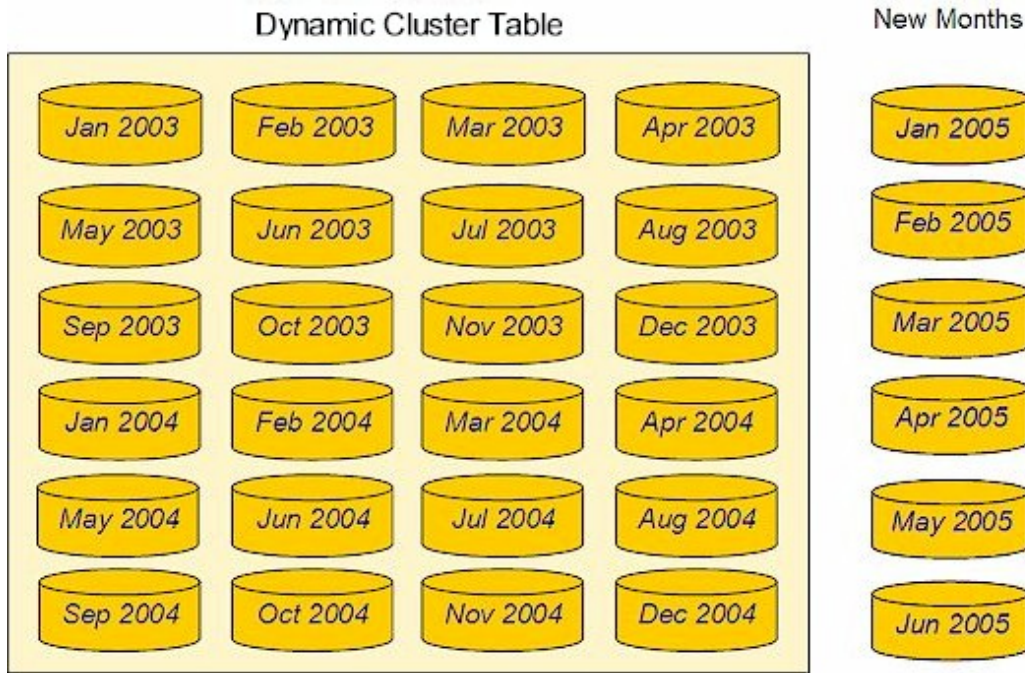
Dynamic Cluster Table Access Control

A user must have SPD Server control access on any member tables that are used in the CLUSTER CREATE or CLUSTER ADD commands. A user must also have SPD Server control access on the dynamic cluster table itself to submit a CLUSTER UNDO command. There is no restriction on table ownership, as long as the user has control access on all member tables. All users that have access to a domain have default control access on tables that were created by the user Anonymous within that domain. ACLs can be defined on a dynamic cluster table after it is created, and the permissions that are specified in the dynamic cluster table ACL are applied when SPD Server accesses the dynamic cluster table. Any individual ACL that is defined on a member table does not apply during the time when the member table is part of a created dynamic cluster table.

Add Tables to a Dynamic Cluster

To add tables to a dynamic cluster table, you must have an existing dynamic cluster table. The SPD Server tables that you want to add to the dynamic cluster table must all be in the same domain as the dynamic cluster table. These tables must use identical table structures (columns and indexes) and compression. However, partition sizes and owners can vary. These requirements ensure the metadata compatibility that is required to add to a dynamic cluster table.

Once the tables to be added are organized, a simple PROC SPDO command is used to add the new tables to an existing dynamic cluster table. In the following graphic, sales tables for the first six months of 2005 are set up to be added to the dynamic cluster table that contains monthly sales transaction data for 2003 and 2004:



The following code shows the PROC SPDO command syntax that is used to add new tables to an existing dynamic cluster table:

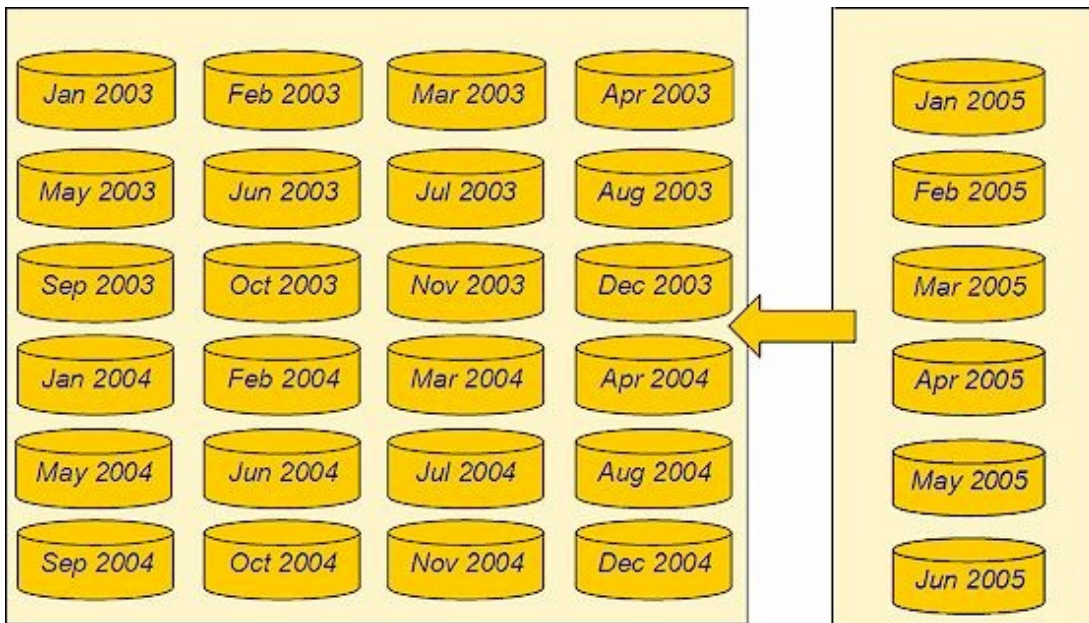
```

PROC SPDO library=domain-name ;
  cluster add Sales_History
    mem=sales200501
    mem=sales200502
    mem=sales200503
    mem=sales200504
    mem=sales200505
    mem=sales200506 ;
quit ;

```

PROC SPDO uses a LIBRARY statement to identify the domain that contains the existing dynamic cluster table that you want to add to. The **cluster add** syntax specifies the name of the dynamic cluster table that you want to add to (Sales_History).

The **mem=** syntax identifies the members that form the table to be added to the existing dynamic cluster table. In the following graphic, six tables that include monthly sales transactions for the first half of 2005 are set up to be added to the existing dynamic cluster table of 2003 and 2004 sales transactions data:




See the "[Dynamic Cluster Table Examples](#)" section for a more extensive code example of adding to a dynamic cluster table.

Undo Dynamic Cluster Tables

To undo a dynamic cluster table, you must have an existing dynamic cluster table. Undoing the dynamic cluster table simply reverts the table back to unbound SPD Server tables. Undoing a dynamic cluster table is required to remove a specific member table from a dynamic cluster table, to add data to a specific member table in the dynamic cluster table, or to completely refresh a specific member table that belongs to the dynamic cluster table.

The following graphic represents a dynamic cluster table with 24 members. Each member contains monthly sales transactions for the years 2003 and 2004:

Dynamic Cluster Table



PROC SPDO is used to undo the existing dynamic cluster table.

The following code shows the PROC SPDO command syntax that is used to undo an existing dynamic cluster table:

```
PROC SPDO library=  
domain-name ;  
cluster undo Sales_History ;  
quit ;
```

PROC SPDO uses a LIBRARY statement to identify the domain that contains the existing dynamic cluster table that you want to undo. The **cluster undo** syntax specifies the name of the dynamic cluster table that you want to undo (Sales_History).

The following graphic represents the previous dynamic cluster table, now unbound.




See the "[Dynamic Cluster Table Examples](#)" section for a more extensive code example of undoing a dynamic cluster table and then refreshing it.

[Refresh Dynamic Cluster Tables](#)

To refresh a dynamic cluster table, you perform the same actions that are required to undo a dynamic cluster table. Then, you recreate the dynamic cluster table after you add a member table or change an existing member table. An example of refreshing an SPD Server dynamic cluster table is updating on a monthly basis a dynamic cluster table whose members are the 24 previous months of sales transaction data.

To refresh a dynamic cluster table, use sequential PROC SPDO commands to UNDO CLUSTER and CREATE CLUSTER with the desired member tables. The dynamic cluster table is first undone. Table changes are made, and then the dynamic cluster table is rebound again. The following example unbinds the sales transactions tables for 2003 and 2004, and then refreshes the dynamic cluster table with sales transactions tables for the first six months of 2005:



See the "[Dynamic Cluster Table Examples](#)" section for a more extensive code example of unbinding a dynamic cluster table and then refreshing it by recreating it with different member tables.

Modify Dynamic Cluster Tables

The PROC SPDO command set for dynamic clusters provides a CLUSTER MODIFY cluster command. The usage syntax for the CLUSTER MODIFY command is

```
CLUSTER MODIFY clustname  
MINMAXVARLIST=(varname1 <varname2 varname3 ...>);
```

The CLUSTER MODIFY command sets a MINMAXVARLIST attribute on one or more variables that belong to an existing dynamic cluster. The variable names that are specified in the CLUSTER MODIFY command must exist in the cluster and the variables must not have a pre-existing MINMAXVARLIST setting. When the SPD Server runs the CLUSTER MODIFY command, the dynamic cluster is unclustered while the variable modifications are made to the individual member tables. The cluster is recreated after the MINMAXVARLIST changes are completed. Control permission and exclusive access to the dynamic cluster is required in order to run the CLUSTER MODIFY command. SPD Server performs a full table scan to initialize the MINMAXVARLIST values in each member table, so the processor time that is required to perform the CLUSTER MODIFY command is directly related to the size of the tables that belong to the cluster. If an error occurs while the CLUSTER MODIFY command is running, the cluster may not be recreated and the user will need to manually recreate the cluster using the CLUSTER CREATE command.

Dynamic Cluster BY Clause Optimization

When you use SPD Server dynamic clusters, you can create huge data sets. If the huge data sets need further manipulation by some SAS job, it might be better to sort them for more efficient processing. Traditional processing of huge data sets can overuse or overwhelm available resources. The resulting lack of available run-time or processor resources can prohibit you from running full-table scans and manipulating table rows, which are required to sort huge data sets for subsequent processing.

SPD Server provides dynamic cluster BY clause optimization to reduce the need for a large amount of processor resources when evaluating BY clauses. The dynamic cluster BY clause optimization uses SPD Server to join individually created SPD Server member data sets so that the data sets appear to be single data set, while still keeping the individual member data sets intact. The dynamic cluster BY clause optimization uses the SORTEDBY metadata of the member data sets to bypass most of the sorting that is required to perform the implicit BY clause ordering. With the SORTEDBY metadata of each member, SPD Server merges the member data sets in the dynamic cluster by using each member data set's order. No additional SPD Server work-space is required, and the ordered data set records are returned with minimum delay since member sorting is eliminated.

To use dynamic cluster BY clause optimization, you need to build the dynamic cluster table a certain way. All of the member tables in your dynamic cluster table need to be sorted by the same columns that you need to use in the BY clause. When you build your dynamic cluster table from member tables that are presorted by your BY clause columns, your dynamic cluster table can use the BY clause optimization.

When a BY clause is run that matches the SORTEDBY column order of the dynamic cluster table member tables, SPD Server performs the BY clause without using sort work-space or experiencing first-record latency. SPD Server uses the presorted member tables to perform an instantaneous interleave. By using the presorted member tables, the dynamic cluster BY clause optimization enables you to perform operations on huge data sets that would be impossible to handle otherwise.

For example, suppose that you have a system that has sufficient CPU, memory, and work-space resources to sort a 50-GB data set in a reasonable amount of time. However, suppose this system accumulates 50 GB of new data every month, so that after 12 months, the data sets require 600 GB of storage. The system cannot handle sorting 600 GB of data to process queries that are based on the previous 12-month period. If you use SPD Server to create a dynamic cluster table from the 12 50-GB member tables, you can store each rolling month of data in a SPD Server member table, and then sort it like the other dynamic cluster table member tables, and then add the new member table to the 600-GB dynamic cluster table. Now you can use the dynamic cluster BY clause optimization to run SAS steps that use BY clauses on the 600-GB cluster. For example, you can run a DATA step MERGE statement that uses the dynamic cluster table as the master source for the MERGE statement. The BY clause from the MERGE statement triggers the dynamic cluster BY clause optimization. As a result, the operation completes in the time that it takes to interleave the individual member tables, using no SPD Server work-space and without experiencing any implicit BY sort delays.

Dynamic cluster BY clause optimization allows the BY optimization to be combined with certain WHERE clauses on dynamic cluster tables. For the WHERE clause optimization to work, SPD Server must be able to determine whether the WHERE clause is trivially true or trivially false for each member table in the dynamic cluster table. To be trivially true, a WHERE clause must find the clause condition true for every row in the member table. To be trivially false, a WHERE clause must find the clause condition false for every row in the member table.

SPD Server keeps metadata about indexed values in dynamic cluster table member tables, and if the WHERE clause criteria can be determined as true or false based on the dynamic cluster table's member table metadata, the WHERE clause optimization is possible on a member-by-member basis for the entire dynamic cluster table. Suppose that member tables of a dynamic cluster table all have an index on the column QUARTER (1=Jan-Mar, 2=Apr-Jun, 3=Jul-Sep, 4=Oct-Dec). Suppose that you need to run a DATA step MERGE statement that uses the expression WHERE QUARTER=2. Because the QUARTER column is indexed in all of the member tables, SPD Server uses the BY clause optimization to determine that the WHERE clause is trivially true. SPD Server then evaluates the expression only on the member tables for April, May, and June without using any SPD Server work-space. When the WHERE clause can be determined as trivially true or trivially false for each member table of the dynamic cluster table in advance, the BY clause optimization performs the BY processing only on the appropriate member tables.

The dynamic cluster BY clause optimization is triggered when member tables all have an applicable SORTEDBY ordering for the BY clause that is asserted. When the SORTEDBY ordering is strong (validated), SPD Server does not perform checks to verify the order of BY variables that are returned from the member table. When the SORTEDBY ordering is weak (such as from a SORTEDBY assertion that was a data set option), additional checking is performed to verify the order of BY variables that are returned from the member table. If an invalid BY variable order is detected, SPD Server terminates the BY clause and displays the following error message:

```
ERROR: Clustered BY member violates weak  
sort order during merge.
```

[Dynamic Cluster BY Clause Optimization Example](#)

Consider a database of medical patient insurance claims, with quarterly claims data sets that are named ClaimsQ1, ClaimsQ2,

ClaimsQ3, and ClaimsQ4. Each quarterly claims table is sorted by columns that are named PatID (for Patient ID) and ClaimID (for Claim ID). The member tables are combined into a dynamic cluster table that is named ClaimsAll. The following example shows the code:

```
DATA SPDS.ClaimsQ1;
...
run;

DATA SPDS.ClaimsQ2;
...
run;

PROC SORT DATA=SPDS.ClaimsQ1;
  BY PatID ClaimID;
run;

PROC SORT DATA=SPDS.ClaimsQ2;
  BY PatID ClaimID;
run;

PROC SPDO LIB=SPDS;
  create cluster ClaimsAll;
quit;
```

Consider the DATA step MERGE statement to be submitted to the ClaimsAll dynamic cluster table:

```
DATA SPDS.ToAdd SPDS.ToUpdate;
MERGE SPDS.NewOnes(IN=NEW1)
      SPDS.ClaimsAll(IN=OLD1);
BY PatID ClaimID;

SELECT;
WHEN(NEW1 and OLD1)
  DO;
  OUTPUT SPDS.ToUpdate;
  end;
WHEN(NEW1 and not OLD1)
  DO;
  OUTPUT SPDS.ToAdd;
  end;
run;
```

If ClaimsAll were not a dynamic cluster table, the DATA step MERGE statement would create an implicit sort from the BY clause on the respective SPD Server data sets. However, ClaimsAll is a dynamic cluster table with member tables that are presorted. As a result, the dynamic cluster BY clause optimization uses BY clause processing to merge the sorted member tables instantaneously without using any SPD Server work-space or experiencing delays. The previous example merges the transaction data named NewOnes into new rows that will be appended to the data for the next quarter.

Consider that the member data sets ClaimsQ1 and ClaimsQ2 are indexed on the column Claim_Date:

```
DATA SPDS.RepClaims;
  SET SPDS.ClaimsAll;
  WHERE Claim_Date BETWEEN '01JAN2007' and '31MAR2007';
  BY PatID ClaimID;
run;
```

The WHERE clause determines whether each member table is true or false for each quarter. The WHERE clause is trivially true for the data set ClaimsQ1 because the WHERE clause is true for all dates in the first quarter. The WHERE clause is

trivially false for the data set ClaimsQ2 because the WHERE clause is false for all dates in the second quarter. The BY clause optimization determines that the member table ClaimsQ1 will be processed because the WHERE clause is true for **all** of the rows of the ClaimsQ1 table. The BY clause optimization skips the member table ClaimsQ2 because the WHERE clause is false for **all** of the rows of the ClaimsQ2 table.

Suppose that the Claim_Date range is changed in the WHERE clause:

```
DATA SPDS.RepClaims;  
SET SPDS.ClaimsAll;  
WHERE Claim_Date BETWEEN '05JAN2007' and '28JUN2007';  
BY PatID ClaimID;  
run;
```

When the new WHERE clause is evaluated, it is not trivially true for member tables ClaimsQ1 or Claims Q2. The WHERE clause is not trivially false for member tables ClaimsQ1 or Claims Q2 either. The WHERE clause calls dates that exist in portions of the member table ClaimsQ1, and it calls dates that exist in portions of the member table ClaimsQ2. The dates in the WHERE clause do not match **all** of the dates that exist in the member table ClaimsQ1, or **all** of the dates that exist in the member table ClaimsQ2. The dates in the WHERE clause are not totally exclusive of the dates that exist in the member tables ClaimsQ1 or ClaimsQ2. As a result, BY clause optimization will not be used when SPD Server runs the code.

Member Table Requirements for Creating Dynamic Cluster Tables

When you create a dynamic cluster table, all of the member tables must have matching table, variable, and index attributes. If there are attribute mismatches, the dynamic cluster table creation fails, and SPD Server displays the following error message:

```
ERROR: Member table not compatible with other  
cluster members. Compare CONTENTS.
```

A more detailed error message is written to the SPD Server log. The SPD Server log lists which attribute is mismatched in the member tables. The following lists specify the member table attributes that must match for SPD Server to successfully create a dynamic cluster table.

- [Table Attributes](#)
- [Variable Attributes](#)
- [Index Attributes](#)

Table Attributes

The following table attributes must match in all member tables to successfully create a dynamic cluster table:

IDXSEGSIZE

index segment size

OBSLEN

observation record length

NVAR

number of columns

NINDEXES

number of indexes

DSORG

data set organization

SEMTYPE

data set semantic type

DSTYPE

SAS data set type

LOCALE

creation locale

LANG

data set language tag

LTYPE

data set language type tag

FLAGS

compressed data set
encrypted data set
backup data set
NLS variables in data set
minmaxvarlist variables in data set
SAS encryption password in data set

SASPW

SAS encryption password

DS_ROLE

data set option for ROLE

ENCODING_CEI

encoding CEI for NLS (for compressed tables)

DISKCOMP

compression algorithm

IOBLOCKSIZE

I/O block size

IOBLOCKFACTOR

I/O block factor

Variable Attributes

The following variable attributes must match in all member tables to successfully create a dynamic cluster table:

NAME

variable name

LABEL

variable label

NFORM

variable format

NIFORM

variable informat

NPOS
variable offset in record

NVARO
variable number in record

NLNG
variable length

NPREC
variable precision

FLAGS
NLS encoding supported
minmaxvarlist variable

NFL
format length

NFD
format decimal places

NIFL
informat length

NIFD
informat precision

NSCALE
scale for fixed-point decimal

NTATTR
variable type attributes

TYPE
variable type

SUBTYPE
variable subtype

SORT
variable sorted status

NTYPE2
variable extended type code

Index Attributes

The following index attributes must match in all member tables to successfully create a dynamic cluster table:

NAME
index name

TYPE
index type

KEYFLAGS

unique index
nomiss index

LENGTH

index length

NVAR

number of variables in index

NVAR0

variable number in index

Querying and Reading Member Tables in a Dynamic Cluster

Dynamic clusters can be read using the MEMNUM= table option. The MEMNUM= option enables you to perform query or read operations on a single member table that belongs to a dynamic cluster. When you use the MEMNUM= option, SPD Server opens only the specified member table, instead of opening all of the member tables that belong to the cluster. You can determine the member number of a table in the cluster by issuing a CLUSTER LIST statement or a PROC CONTENTS command on the cluster. The SPD Server CLUSTER LIST or PROC CONTENTS command output lists the member tables of the cluster in numbered order.

You can specify an integer value n as an argument for the MEMNUM= table option to select the n th member of the table, or you can use the argument LASTCLUSTERMEMBER. When you use the LASTCLUSTERMEMBER argument with MEMNUM=, SPD Server selects the last member of the dynamic cluster table, without needing to count the members to determine the number (n) of the last member.

The following example uses the MEMNUM= table option to query against the member table sales200504 that belongs to the dynamic cluster table sales_history:

```
PROC SPDO library=&domain ;
  CLUSTER CREATE sales_history
    mem=sales200501
    mem=sales200502
    mem=sales200503
    mem=sales200504
    mem=sales200505
    mem=sales200506
    maxslot=12 ;
quit ;

PROC PRINT data=&domain..sales_history
  (MEMNUM=4);
  WHERE salesdate = 30Apr2005;
run;
```

To use the MEMNUM= table option to query the last member table in the dynamic cluster table sales200506, the query would be:

```
PROC SPDO library=&domain ;
  CLUSTER CREATE sales_history
    mem=sales200501
    mem=sales200502
    mem=sales200503
    mem=sales200504
    mem=sales200505
    mem=sales200506
    maxslot=12 ;
quit ;
```



```
PROC PRINT data=&domain..sales_history
(MEMNUM=LASTCLUSTERMEMBER);
WHERE salesdate = 15Jun2005;
run;
```

Unsupported Features in Dynamic Cluster Tables

Because of differences in the load and read structures for dynamic cluster tables, some standard features that are available in SAS tables and SPD Server tables are currently not supported in SPD Server 4.4. These features are:

- You cannot append data to a dynamic cluster table. To append data to a dynamic cluster table, the table must be unclustered, the data is appended to the individual unclustered files, and then the individual unclustered files must be reclustered.
- Record-level locking is not allowed.
- The SPD Server backup/restore utility is not available.
- Copying data with PROC COPY or PROC SQL is not supported.

If a task for a dynamic cluster table requires one of these features, you should undo the dynamic cluster table and create standard SPD Server tables.

Dynamic Cluster Table Examples

The following four examples show all of the fundamental operations that are required to use dynamic cluster tables:

- [Create a Dynamic Cluster Table Example](#)
- [Add Tables to a Dynamic Cluster Example](#)
- [Undo Dynamic Cluster Table Example](#)
- [Refresh Dynamic Cluster Table Example](#)

Create a Dynamic Cluster Table Example

The following example creates a dynamic cluster table named Sales_History. The first part of the example generates dummy transaction data that is used in the rest of the example.

The example uses SPD Server tables from the domain bmwcycle. Twelve individual SPD Server tables for monthly motorcycle sales during 2004 are bound into the dynamic cluster table named Sales_History. Tables are created for the first six months of motorcycle sales during 2005:

```
/* declare macro variables that will be used to */
/* generate dummy transaction data */

%macro var (varout,dist,card,seed,peak) ;
  %put &dist &card &seed ;
  %local var1 ;

  if upcase("&dist") = 'RANUNI'
  then do ;
    &varout = int(ranuni(&seed)*&card)+1;
  end ;
```

```

else
if upcase("&dist") = 'RANTRI'
then do ;
    *%let vartri = %substr("&dist",5,2)&card ;
    *&varout = int(rantri(&seed,&peak)*&card)+1;
    &varout = int(rantri(&seed,&peak)*&card)+1;
end ;
%mend ;

%macro linkvar (varin,varout,devisor) ;
    &varout = int(&varin/&devisor) ;
%mend ;

/* declare main vars */

%let domain=bmwcycle ;
%let host=kaboom ;
%let port=5200 ;
%let spdssize=256M ;
%let spdsiasy=YES ;

libname &domain sasspds "&domain"
    server=&host.&port
    user='anonymous'
    ip=YES ;

/* generate monthly sales data tables for */
/* 2004 and the first six months of 2005 */

data
    &domain..sales200401
    &domain..sales200402
    &domain..sales200403
    &domain..sales200404
    &domain..sales200405
    &domain..sales200406
    &domain..sales200407
    &domain..sales200408
    &domain..sales200409
    &domain..sales200410
    &domain..sales200411
    &domain..sales200412
    &domain..sales200501
    &domain..sales200502
    &domain..sales200503
    &domain..sales200504
    &domain..sales200505
    &domain..sales200506
    ;

drop seed bump1 bump2 random_dist ;

seed = int(time()) ;

/* format the dummy transaction data */

```

```

format trandate shipdate paidddate yymmdd10. ;

put seed ;
do transact = 1 to 5000 ;
  %var (customer,ranuni,100000,seed,1) ;

  %linkvar (customer,zipcode,10) ;
  %linkvar (customer,agent,20) ;
  %linkvar (customer,mktseg,10000) ;
  %linkvar (agent,state,100) ;
  %linkvar (agent,branch,25) ;
  %linkvar (state,region,10) ;

  %var (item_number,ranuni,15000,seed,1) ;

  %var (trandate,ranuni,577,seed,1) ;
  trandate = trandate + 16071 ;

  %var (bump1,ranuni,20,seed,.1) ;
  shipdate = trandate + bump1 ;

  %var (bump2,rantri,30,seed,.5) ;
  paidddate = trandate + bump2 ;

  %var (units,ranuni,100,seed,1) ;
  %var (trantype,ranuni,10,seed,1) ;
  %var (amount,rantri,50,seed,.5) ;
  amount = amount + 25 ;

  random_dist = ranuni ('03feb2005'd) ;

  /* sort the dummy transaction data into */
  /* monthly sales data tables          */

  if '01jan2004'd <= trandate <= '31jan2004'd
    then output &domain..sales200401 ;

  else if '01feb2004'd <= trandate <= '28feb2004'd
    then output &domain..sales200402 ;

  else if '01mar2004'd <= trandate <= '31mar2004'd
    then output &domain..sales200403 ;

  else if '01apr2004'd <= trandate <= '30apr2004'd
    then output &domain..sales200404 ;

  else if '01may2004'd <= trandate <= '31may2004'd
    then output &domain..sales200405 ;

  else if '01jun2004'd <= trandate <= '30jun2004'd
    then output &domain..sales200406 ;

  else if '01jul2004'd <= trandate <= '31jul2004'd
    then output &domain..sales200407 ;

  else if '01aug2004'd <= trandate <= '31aug2004'd
    then output &domain..sales200408 ;

```

```

else if '01sep2004'd <= trandate <= '30sep2004'd
  then output &domain..sales200409 ;

else if '01oct2004'd <= trandate <= '31oct2004'd
  then output &domain..sales200410 ;

else if '01nov2004'd <= trandate <= '30nov2004'd
  then output &domain..sales200411 ;

else if '01dec2004'd <= trandate <= '31dec2004'd
  then output &domain..sales200412 ;

else if '01jan2005'd <= trandate <= '31jan2005'd
  then output &domain..sales200501 ;

else if '01feb2005'd <= trandate <= '28feb2005'd
  then output &domain..sales200502 ;

else if '01mar2005'd <= trandate <= '31mar2005'd
  then output &domain..sales200503 ;

else if '01apr2005'd <= trandate <= '30apr2005'd
  then output &domain..sales200504 ;

else if '01may2005'd <= trandate <= '31may2005'd
  then output &domain..sales200505 ;

else if '01jun2005'd <= trandate <= '31jun2005'd
  then output &domain..sales200506 ;
end ;
run ;

/* index the transaction data in the */
/* monthly sales data tables      */

%macro indexit (yrmth) ;
  PROC DATASETS library=&domain nolist ;
    modify sales&yrmth ;
    index create transact customer agent state branch trandate ;
  quit ;
%mend ;

%let spdsiasy=YES ;

%indexit (200401) ;
%indexit (200402) ;
%indexit (200403) ;
%indexit (200404) ;
%indexit (200405) ;
%indexit (200406) ;
%indexit (200407) ;
%indexit (200408) ;
%indexit (200409) ;
%indexit (200410) ;
%indexit (200411) ;
%indexit (200412) ;
%indexit (200501) ;
%indexit (200502) ;

```

```

%indexit (200503) ;
%indexit (200504) ;
%indexit (200505) ;
%indexit (200506) ;

/* Use PROC SPDO to create the dynamic cluster */
/* table sales_history */

PROC SPDO library=&domain ;
  cluster create sales_history
    mem=sales200401
    mem=sales200402
    mem=sales200403
    mem=sales200404
    mem=sales200405
    mem=sales200406
    mem=sales200407
    mem=sales200408
    mem=sales200409
    mem=sales200410
    mem=sales200411
    mem=sales200412
  maxslot=36 ;
quit ;

```

Add Tables to a Dynamic Cluster Example

The following example adds member tables to the dynamic cluster table named Sales_History. The Sales_History table currently contains 12 members. Each member is an SPD Server table that contains monthly sales data. This example augments the 12 member tables for 2004 with 6 new member tables that contain sales data for January through June of 2005:

```

/* declare main vars */

%let domain=bmwcycle ;
%let host=kaboom ;
%let port=5200 ;
%let spdssize=256M ;
%let spdsiasy=YES ;

libname &domain sasspds "&domain"
  server=&host.&port
  user='anonymous'
  ip=YES ;

/* Use PROC SPDO to add member tables to */
/* the dynamic cluster table sales_history */

PROC SPDO library=&domain ;
  cluster add sales_history
    mem=sales200501
    mem=sales200502
    mem=sales200503
    mem=sales200504
    mem=sales200505

```

```

        mem=sales200506 ;
quit ;

/* Verify the presence of the added tables */

PROC CONTENTS data=&domain..sales_history ;
run ;

```

Undo Dynamic Cluster Table Example

The undo example is included as part of the following refresh example.

Refresh Dynamic Cluster Table Example

Refreshing SPD Server dynamic cluster tables is a combination of two tasks, UNDO CLUSTER and CREATE CLUSTER. The UNDO CLUSTER command unbinds an existing dynamic cluster table. The CREATE CLUSTER command rebinds the dynamic cluster table with updated member tables. Therefore, the following example shows both the UNDO CLUSTER and CREATE CLUSTER commands with SPD Server dynamic cluster tables.

The following example refreshes the dynamic cluster table named Sales_History. The Sales_History table received additional member tables in the previous example. The 18-member dynamic cluster table Sales_History is unbound. The 12 member tables that contain 2004 sales data are deleted when the dynamic cluster table Sales_History is recreated with only the six member tables that contain 2005 sales data. The combined actions refresh the contents of the dynamic cluster table Sales_History.

```

/* declare main vars */

%let domain=bmwcycle ;
%let host=kaboom ;
%let port=5200 ;
%let spdssize=256M ;
%let spdsiasy=YES ;

libname &domain sasspds "&domain"
        server=&host..&port
        user='anonymous'
        IP=YES ;

/* Use PROC SPDO to undo the existing dynamic */
/* cluster table Sales_History, then rebind */
/* it with members from months in 2005 only */

PROC SPDO library=&domain ;
cluster undo sales_history ;
cluster create sales_history
        mem=sales200501
        mem=sales200502
        mem=sales200503
        mem=sales200504
        mem=sales200505
        mem=sales200506
        maxslot=36 ;
quit ;

```

```
/* Verify the contents of the refreshed dynamic */  
/* cluster table sales_history */  
  
PROC CONTENTS data=&domain..sales_history ;  
run ;
```

Copyright © 2008 by SAS Institute Inc., Cary, NC, USA. All rights reserved.

SAS Scalable Performance Data (SPD) Server SQL Features

- [SPD Server SQL Planner](#)
 - [Connecting to the SPD Server SQL Engine](#)
 - [SPD Server SQL Planner Options](#)
 - [Specify SQL Options using Explicit Pass-Through Code](#)
 - [Specify SQL Options using Implicit Pass-Through Code](#)
 - [Important SPD Server SQL Planner Options](#)
 - [Parallel Join Facility](#)
 - [Parallel Join Methods](#)
 - [Parallel Joins with Group-By](#)
 - [Parallel Join SQL Options](#)
 - [Parallel Join Example 1](#)
 - [Parallel Join Example 2](#)
 - [Parallel Join Example 3](#)
 - [Parallel Group-By Facility](#)
 - [Enhanced Group-By Functions](#)
 - [Table Aliases Supported](#)
 - [Nested Queries Meet Group-By Syntax Requirements](#)
 - [Formatted Parallel Group Select](#)
 - [Parallel Group-By SQL Options](#)
 - [SPD Server STARJOIN Facility](#)
 - [STARJOIN Options](#)
 - [SPD Server Index Scan](#)
 - [Optimizing Correlated Queries](#)
 - [Correlated Query Options](#)
 - [Materialized Views](#)
 - [SPD Server SQL Extensions](#)
 - [Differences between SAS SQL and SPD Server SQL](#)
-

[SPD Server SQL Planner](#)

SPD Server includes SQL Planner optimizations. SQL Planner optimizations improve the performance of the more frequent query types that used in data mining solutions such as Enterprise Marketing Automation. A key enhancement to the SPD Server SQL Planner is optimizing correlated queries through the use of query rewrite techniques. Correlated queries are common in business and analytic intelligence data mining. Another significant enhancement is the tighter integration of the Parallel Group-By technology in the planner. The tighter integration adds performance benefits to nested Group-By syntax.

[Connecting to the SPD Server SQL Engine](#)

- [Implicit Pass-Through Connection](#)
- [Explicit Pass-Through Connection](#)
 - [LIBNAME Syntax to Specify a LIBREF](#)
 - [LIBREF Statements](#)
 - [LIBREF Clauses](#)
 - [LIBREF Examples](#)

[Implicit Pass-Through Connection](#)

You can use an implicit pass-through connection to pass implicit SQL statements to the SPD Server SQL Engine. When you use an implicit pass-through connection, the SAS SQL planner parses SQL statements to determine which, if any, portions can be passed to the SPD Server SQL Engine. In order for a submitted SQL statement to take advantage of implicit pass-through SQL, the tables that are referenced in the SQL statement must be SPD Server tables, and the SPD Server SQL engine must be able to successfully parse the submitted SQL statement.

An example of an SPD Server implicit pass-through connection is available in the Help section in this document on how to [specify SQL options using implicit pass-through code](#).

Explicit Pass-Through Connection

You can use an explicit pass-through connection to pass explicit SQL statements to the SPD Server SQL Engine. When you use an explicit pass-through connection, you decide explicitly which SQL statements are passed to the SPD Server SQL Engine. The explicit pass-through connection passes the entire SQL statement as written to the SPD Server SQL Engine, which parses and plans the SQL statement. All tables that are referenced in the SQL statement must be SPD Server tables or an error will occur.

An example of an SPD Server implicit pass-through connection is available in the Help section in this document on how to [specify SQL options using explicit pass-through code](#).

LIBNAME Syntax to Specify a LIBREF

Below is a LIBNAME statement that associates a libref, the SASSPDS engine, and an SPD Server domain.

```
LIBNAME libref  
SASSPDS <'SAS-data-library'> <SPD Server-options>;
```

Use the following arguments:

libref

a name that is up to eight characters long and that conforms to the rules for SAS names.

SASSPDS

the name of the SPD Server engine.

'SAS-data-library'

the logical LIBNAME domain name for an SPD Server data library on the host machine. The name server resolves the domain name into the physical path for the library.

SPD Server-options

one or more SPD Server options.

LIBREF Statements

Whenever you issue a [CONNECT statement](#) to an SPD Server SQL server with the DBQ option, by default you define a primary LIBNAME domain. The software uses the primary domain to resolve table references in SQL statements executed for that connection.

You can also use the LIBREF statement to assign secondary LIBNAME domains for the SPD Server SQL Server. The additional LIBREF statements assign explicit LIBNAME domains, allowing the software to specify two-part table names for SQL statements executed for the connection.

```
PROC SQL;  
execute(libref librefname  
  <enginename>  
  engopt= ' ')  
by sasspds;
```

LIBREF Clauses

- [The ENENAME Clause](#)
- [The ENGOPT Clause](#)

The ENENAME Clause

Specifies the name of an alternate SAS I/O engine to service the LIBREF's access to data. If you do not specify an alternate SAS I/O engine, the default is **spdseng**, which accesses SPD Server tables.

The ENGOPT Clause

Specifies options that configure the LIBREF to access a specific data source or storage domain. Use single or double quotes around the clause. (If you have nested quotes within a clause, alternate between single and double quoted expressions.) The available options depend on the current value of the ENGNAME option. For the default **spdseng**, you can specify any SPD Server CONNECT or LIBNAME engine option with the exception of **prompt**, **newpasswd**, and **chnypass**. Use the same keyword/value syntax required by the [CONNECT statement](#).

Note: If you specify the SAS I/O engine **spdseng** and use explicit options in your CONNECT statement, these options become default ENGOPT clause options. Explicit options can also be specified using the ENGOPT clause. Explicit options specified in an ENGOPT clause will override default values or declarations made in previous CONNECT statements.

LIBREF Examples

- [LIBREF for Another Domain but the Same CONNECT Statement User](#)
- [LIBREF to Same Domain but Different CONNECT Statement User](#)
- [Secondary LIBREF Using a Different Host](#)

LIBREF for Another Domain but the Same CONNECT Statement User

In this example the client connects to the SPD Server SQL server using the engine **sasspds**. The domain is "mydomain," the server machine is called "namesvrID," and the port number is "namesvrPortNum". The execute statement assigns the LIBREF "cookie" to another domain, "dough." After the libref is executed, the user issuing the connect statement can now access either the default domain "mydomain" or the secondary domain "dough."

```
PROC SQL;
connect to sasspds
  (dbq='mydomain'
   host='namesvrID'
   serv='namesvrPortNum'
   user='neraksr'
   passwd='siuya');

execute(libref cookie
  engopt='
  dbq="dough"')
by sasspds;
```

In the example above, the LIBREF is "cookie," and the secondary domain named is "dough." The intent of the example is to show how the CONNECT and LIBREF statements work in conjunction to access multiple domains for the same user.

LIBREF to Same Domain but Different CONNECT Statement User

This example assigns a LIBREF to the domain specified by the CONNECT statement but for another user (different SPD Server User ID).

```
PROC SQL;
execute(libref samslib
  engopt='
  user="sam"
  passwd="samspwd"')
by sasspds;
```

Secondary LIBREF Using a Different Host

This example assigns a secondary LIBREF to a different host machine.

```
PROC SQL;
execute(libref sam2
engopt='
host="flex"
dbq="samsplace"')
by sasspds;
```

SPD Server SQL Planner Options

The SPD Server SQL Planner provides reset options that you can use to configure the behavior of the SQL Planner and the SPD Server facilities that function through the SQL Planner, such as the SPD Server Parallel Group-By facility, the SPD Server Parallel Join facility, and the SPD Server STARJOIN facility. You can specify SPD Server SQL reset options using either using [explicit pass-through](#) code, or by using [implicit pass-through](#) code.

Specify SQL Options using Explicit Pass-Through Code

The example below shows how to use an `execute(reset <reset-options>)` statement in explicit SPD Server pass-through SQL code to invoke an SQL Planner, Parallel Group-By facility, Parallel Join facility, or STARJOIN facility reset option.

Most SQL Planner reset option usage examples in this document use explicit pass-through code. See the implicit pass-through code example below to see how SQL reset options can be declared using an implicit `%let spdssqlr=` statement instead of an explicit `execute(reset <reset-options>)` statement.

```
/* Explicit Pass-Through SQL Example */
/* to invoke an SQL Reset Option    */

PROC SQL ;

connect to sasspds (
  dbq="domain-name"
  server=<host-name>.<port-number>
  user='username' ) ;

execute(reset <reset-options>)
  by sasspds ;

execute(SQL statements)
  by sasspds ;

disconnect from sasspds ;
quit ;
```

Specify SQL Options using Implicit Pass-Through Code

The example below shows how to use a `%let spdssqlr=<reset-options>` statement in implicit SPD Server pass-through SQL code to invoke an SQL Planner, Parallel Group-By facility, Parallel Join facility, or STARJOIN facility reset option.

Most SQL Planner reset option usage examples in this document use explicit pass-through code. The implicit pass-through code example below shows how SQL reset options can be declared using an implicit `%let spdssqlr=` statement instead of an explicit `execute(reset <reset-options>)` statement.

```
/* Implicit Pass-Through SQL Example */
/* to invoke an SQL Reset Option    */

%let spdssqlr=<reset-options> ;
```

```
PROC SQL ;
SQL statements ;

quit ;
```

Important SPD Server SQL Planner Options

The following reset options belong directly to the SPD Server SQL Planner.

- [_Method](#)
- [EXEC/NOEXEC](#)
- [MAGIC](#)
- [INDEXSELECTIVITY=](#)
- [OUTRSRTJNDX/NOOUTRSRTJNDX](#)
- [INOBS](#)
- [OUTOBS](#)
- [SASVIEW/NOSASVIEW](#)
- [UNDO_POLICY=](#)
- [BUFFERSIZE=](#)
- [PRINTLOG/NOPRINTLOG](#)

More detailed information about the available SQL reset options for the SPD Server SQL Parallel Join, Parallel Group-By, STARJOIN, and Correlated Query facilities can be found in this document as follows:

- [Parallel Join Facility Reset Options](#)
- [Parallel Group-By Facility Reset Options](#)
- [STARJOIN Facility Reset Options](#)
- [Correlated Query Facility Options](#)

_Method

The SQL **_method** option is one of the most important reset options. The **_method** reset option provides a method tree in the output that shows how the SQL was executed.

The following methods are displayed in the SQL **_method** tree:

sqxcrta
Create table as Select.

sqxslct
Select rows from table.

sqxjsl
Step Loop Join (Cartesian Join).

sqxjm
Merge Join execution.

sqxjndx
Index Join execution.

sqxjhsh
Hash Join execution.

sqxsort
Sort table or rows.

sqxsrc

Read rows from source.

sqxfil

Filter rows from table.

sqxsumg

Summary Statistics (with GROUP BY).

sqxsumn

Summary Statistics (not grouped).

sqxuniq

Distinct rows only.

sqxstj

STARJOIN

sqxxpqb

Parallel Group-By

sqxxpjn

Parallel Join with Group-By. The SAS log will display the name of the parallel join method that was used.

sqxpjl

Parallel Join without Group-By

Reading the Method Tree

A method tree is produced in your output when the **_method** reset option is specified for the SQL Planner. The SQL Planner method tree is read from bottom row to top row. Below is an example that shows how to interpret the method tree by substituting the type of method that was used in each step.

```
PROC SQL ;
create table tbl1 as
select *
  from path1.dansjunk1 a,
      path1.dansjunk2 b,
      path1.dansjunk3 c
 where a.i = b.i
      and a.i = c.i ;
quit ;
```

Here is the example Method Tree that was printed:

```
SPDS_NOTE: SQL execution methods chosen are:
<0x00000001006BBD78>    sqxslct
<0x00000001006BBBF8>    sqxjm
<0x00000001006BBE38>    sqxsort
<0x0000000100691058>    sqxsrc
<0x0000000100667280>    sqxjm
<0x0000000100666C50>    sqxsort
<0x0000000100690BD8>    sqxsrc
<0x00000001006AE600>    sqxsort
<0x0000000100694748>    sqxsrc
```

Reading from bottom to top, you can review the sequence of methods that were invoked.

```
SPDS_NOTE: SQL execution methods chosen are:
<0x00000001006BBD78>    step-9
<0x00000001006BBBF8>    step-8
<0x00000001006BBE38>    step7
<0x0000000100691058>    step-6
```

<0x0000000100667280>	step-5
<0x0000000100666C50>	step-4
<0x0000000100690BD8>	step-3
<0x00000001006AE600>	step-2
<0x0000000100694748>	step-1

In step 1, **sqxsrc** reads rows from the source. In step 2, **sqxsort** sorts the table rows. Then in steps 3 and 4, more rows are read and sorted. In step 5, the tables are joined by **sqxjm**, and so on.

EXEC/NOEXEC

You use the SPD Server SQL Planner **EXEC/NOEXEC** option to turn SPD Server SQL execution on or off.

Usage:

```

/* This explicit Pass-Through SQL */
/* prints the method tree without */
/* executing the SQL code.        */

PROC SQL ;
connect to sasspds
  (dbq="domain"
   server=<host-name>.<port-number>
   user='username') ;

execute (reset _method noexec)
  by sasspds ; /* turns SQL exec off */

execute (SQL statements)
  by sasspds ;

disconnect from sasspds ;
quit ;

```

MAGIC

You use the SPD Server SQL Planner **MAGIC** reset option that controls how the SPD Server SQL planner executes join statements. The Magic option has three settings, 101, 102, and 103.

Usage:

```

execute(reset magic=<101/102/103>)
  by sasspds ;

```

MAGIC=101

SPD Server performs sequential loop joins. Sequential loop joins are brute force joins that match every row from the first table to every row of the second table.

MAGIC=102

SPD Server performs sort merge joins. Sort merge joins force a sort on all tables that are involved in the join.

MAGIC=103

SPD Server performs hash joins. Hash joins require SPD Server to create a memory table in order to perform the join. The size of the memory table is limited based on memory available.

INDEXSELECTIVITY=

The **INDEXSELECTIVITY=** option allows you to tune SPD Server join methods, depending on the proportion of the tables to be joined. The **indexselectivity** property can have values of 0.0 or 1.0.

Usage:

```
execute(reset indexselectivity=<1.0/0.0>)
  by sasspds ;
```

INDEXSELECTIVITY=0.0 is the default SPD Server setting and uses index join methods.

INDEXSELECTIVITY=1.0 configures the SPD Server SQL planner to use a 15% heuristic rule when it decides to perform a join with an index. If the SPD Server SQL Planner calculates that less than 15% of the total table rows will be selected in a join, the SQL Planner uses an index join method. If the SPD Server SQL Planner determines that greater than 15% of the total table rows will be selected in a join, the SQL Planner uses a sort merge join.

OUTRSRTJNDX/NOOUTRSRTJNDX

Use the **OUTRSRTJNDX/NOOUTRSRTJNDX** option to configure sort behavior for a SPD Server join index. **OUTRSRTJNDX** sorts the outer table for a join index by the join key. This is the default SPD Server setting. **NOOUTRSRTJNDX** does not sort the outer table for a join index.

Usage:

```
/* Disable outer table      */
/* sorting for a join index */
execute(reset nooutrsrtjndx)
  by sasspds ;
```

```
/* Enable outer table      */
/* sorting for a join index */
execute(reset outrsrtjndx)
  by sasspds ;
```

INOBS

Use the **INOBS** option to specify the specific number of observations that you want to read from input tables.

Usage:

```
execute(reset inobs=<n>)
  by sasspds ;
```

where the integer value **<n>** is the desired number of observations.

OUTOBS

Use the **OUTOBS** option to specify the specific number of observations that you want to create or print in your output.

Usage:

```
execute(reset outobs=<n>)
  by sasspds ;
```

where the integer value **<n>** is the desired number of observations.

SASVIEW/NOSASVIEW

Use the **SASVIEW/NOSASVIEW** option to enable or disable SAS PROC SQL views that use an SPD Server LIBNAME. SAS PROC SQL views use a generic transport format to represent numeric values, which SPD Server converts to native numeric values. When extremely large or extremely small numeric values are conveyed in a SAS PROC SQL view to SPD Server, some precision may be lost in extreme values during the SPD Server numeric conversion.

Usage:

```
/* Disable SAS PROC SQL views      */
/* that use an SPD Server LIBNAME */
execute(reset nosasview)
  by sasspds ;

/* Enable SAS PROC SQL views that */
/* use an SPD Server LIBNAME      */
execute(reset sasview)
  by sasspds ;
```

If SAS PROC SQL views are disabled and SPD Server Pass-Through SQL uses a view that was created by PROC SQL, SPD Server rejects the PROC SQL statement and inserts the following error message in the SAS log::

```
SPDS_WARNING: SAS View and SASVIEW Reset Option equals No.
SPDS_ERROR: An error has occurred.
```

If SAS PROC SQL views are enabled and SPD Server Pass-Through SQL uses a view that was created by PROC SQL, SPD Server prints the following note in the SAS log:

```
SPDS_NOTE: SPDS using SAS View in transport mode.
```

UNDO_POLICY=

Use the **UNDO_POLICY** option in SPD Server PROC SQL and RESET statements to configure SPD Server PROC SQL error recovery. When you update or insert rows in a table, you may receive an error message that states that the update or insert operation cannot be performed. The **UNDO_POLICY** option specifies how you want SPD Server to handle rows that were affected by INSERT or UPDATE statements that preceded a processing error.

Usage:

```
/* Do not undo any updates or inserts */
execute(reset undo_policy=none)
  by sasspds ;

/* Permit row inserts and updates to */
/* be done up to the point of error  */
execute(reset undo_policy=required)
  by sasspds ;
```

UNDO_POLICY=NONE

is the default setting for SPD Server. It does not undo any updates or inserts.

UNDO_POLICY=REQUIRED

undoes all row updates or inserts up to the point of error.

UNDO_POLICY=OPTIONAL

Undoes any updates or inserts that it can undo reliably.

If the UNDO policy is not REQUIRED, you will get the following warning message for an insert into the table:

```
WARNING: The SQL option UNDO_POLICY=REQUIRED is not in effect. If an
error is detected when processing this insert statement, that error
will not cause the entire statement to fail.
```

BUFFERSIZE=

The SPD Server query optimizer considers a hash join when an index join is eliminated as a possibility. With a hash join, the smaller table is reconfigured in memory as a hash table. SQL sequentially scans the larger table and row-by-row performs a hash lookup against the small table to form the result set. On a memory-rich system, consider increasing the **BUFFERSIZE=** option to increase the likelihood that a hash join is chosen. The default **BUFFERSIZE=** setting is 64K. You can specify the amount of memory that you want SPD Server to use for hash joins.

Usage:

```
/* Increase buffersize from 64K */
execute(reset buffersize=1048576)
by sasspds ;
```

PRINTLOG/NOPRINTLOG

You use the PRINTLOG/NOPRINTLOG option of the SPD Server SQL Planner to turn the printing of the SQL statement text to the SPD Server log on or off.

Usage:

```
PROC SQL ;
connect to sasspds
  (dbq="domain"
  server=<host-name>.<port-number>
  user='username') ;

/* turn SQL statement printing on */
execute (reset printlog)
by sasspds ;

/* all statements will be printed to SPD Server log */
execute (SQL statements)
by sasspds ;

/* turn SQL statement printing off */
execute (reset noprintlog)
by sasspds ;

disconnect from sasspds ;
quit ;
```

Parallel Join Facility

The Parallel Join facility is a feature of the SPD Server SQL planner that decreases the required processing time when creating a pair-wise join between two SPD Server tables. The processing time savings is created when SPD Server performs the pair-wise join in parallel.

The SQL planner first searches for pairs when SPD Server source tables are to be joined. When a pair is found, the planner checks the join syntax for that pair to determine if it meets all of the requirements for the Parallel Join facility. If the join syntax meets the requirements, the pair of tables will be joined by the Parallel Join facility.

- [Parallel Join Methods](#)
- [Parallel Joins with Group-By](#)
- [Parallel Join SQL Options](#)
- [Parallel Join Example 1](#)
- [Parallel Join Example 2](#)
- [Parallel Join Example 3](#)

[Parallel Join Methods](#)

The SPD Server Parallel Join facility supports two methods, the [parallel sort-merge method](#) and the [parallel range join method](#).

[Parallel Sort-Merge Method](#)

The parallel sort-merge join method first performs a parallel sort to order the data, then merges the sorted tables in parallel. During the merge, the facility concurrently joins multiple rows from one table with the corresponding rows in the other table. You can use the parallel sort-merge join method to execute any join that meets the requirements for parallel join.

The parallel sort-merge method is a good all-around parallel join strategy that requires no intervention from the user. The tables for the sort-merge method do not need to be in the same domain. The performance for the sort-merge method is not affected by the distribution of the data in the sort key columns.

The sort-merge method begins by completely sorting the smaller of the two tables being joined, while also performing concurrent partial parallel sorts on the larger table. If both tables are very large and sufficient resources are not available to do the complete sort on the smaller table, the performance of the parallel sort-merge method can degrade. The parallel sort-merge method is also limited when performing an outer, left, or right join in parallel. Only two concurrent threads can be used when performing parallel outer, left, or right joins. Inner joins are not limited in the parallel sort-merge method and can utilize more than two concurrent threads during parallel operations.

[Parallel Range Join Method](#)

The parallel range join method uses a join index to determine the ranges of rows between the tables that can be joined in parallel. The parallel range join method requires you to create a join index on the columns to be joined in the tables that you want to merge. The join index divides the two tables into a specified number of near-equal parts, or ranges, based on matching values between the join columns. The Parallel Join Facility recognizes the ranges of rows that contain matching values between the join columns, then uses concurrent join threads to join the rows in parallel. The SPD Server parallel sort then sorts the rows within a range.

The parallel range join method can only be performed on tables that are in the same domain. If either of the two tables are updated after the join index is created, the join index must be rebuilt before the parallel range join method can be used. The parallel range join method performs best when the columns of the tables that are being joined are sorted. If the columns are not relatively sorted, then the concurrent join threads can cause processor thrashing. Processor thrashing occurs when unsorted rows in a table require SPD Server to perform increasingly larger table row scans, which can consume processor resources at a high rate during concurrent join operations.

More detailed information on creating join indexes is available in the section in the SPD Server Administrator's Guide documentation on The Hybrid Index Utility Ixutil.

How does the SPD Server Parallel Join facility choose between the sort-merge method and the range join method? If a join index is available for the tables to be joined, the Parallel Join facility will choose the parallel range join method. If a join index does not exist, or if the join index has not been rebuilt since a table was updated, the Parallel Join facility defaults to using the parallel sort-merge method.

[Parallel Joins with Group-By](#)

A powerful feature of the SPD Server Parallel Join facility is its integration with the SPD Server Parallel Group-By

facility. If the result of the parallel join contains a group-by statement, the partial results of the parallel join threads are passed to the Parallel Group-By facility, which performs the group-by operation in parallel. In the following example, SPD Server performs both a parallel join and parallel group-by operation.

```
LIBNAME path1 sasspds .... IP=YES;

PROC SQL;
create table junk as
select a.c, b.d, sum(b.e)
from path1.table1 a,
     path1.table2 b
where a.i = b.i
group by a.d, b.d;
quit;
```

When you use the SPD Server Parallel Join facility, you are not restricted to using the parallel group-by method only on single tables.

Parallel Join SQL Options

SPD Server provides the following Parallel Join reset options:

- [PLLJOIN/NOPLLJOIN](#)
- [CONCURRENCY](#)
- [PLLJMAGIC](#)

PLLJOIN/NOPLLJOIN

The **PLLJOIN/NOPLLJOIN** option enables and disables the SPD Server Parallel Join facility.

Usage:

```
execute(reset noplljoin)
by sasspds ; /* disables Parallel Join */
```

CONCURRENCY

The **CONCURRENCY=<n>** option sets the concurrency level that is used by the SPD Server Parallel Join facility, where the integer *n* specifies the number of levels. In most cases, changing the default SPD Server concurrency setting (half of the available number of processors) is not recommended.

Usage:

```
execute(reset concurrency=4)
by sasspds ; /* enables 4 concurrency levels */
```

PLLJMAGIC

The **PLLJMAGIC** option specifies how SPD server performs parallel joins.

Usage:

```
execute(reset plljmagic=<100/200>)
by sasspds ;
```

PLLJMAGIC=100 forces a parallel range join when the range index is available.

PLLJMAGIC=200 forces a parallel merge join.

Parallel Join Example 1

The first parallel join example is a basic SQL query that creates a pair-wise join of two SPD Server tables, **table1** and **table2**.

```
LIBNAME path1 sasspds .... IP=YES;

PROC SQL;
create table junk as
select *
  from path1.table1 a,
       path1.table2 b
 where a.i = b.i;
quit;
```

Parallel Join Example 2

The next parallel join example is an SQL query that uses more than two SPD Server tables. In this example, the SQL planner performs a parallel join on **table1** and **table2**, and then use a non-parallel method to join the results of the first join and **table3**. A non-parallel join method is used for the second join, because the criteria for a parallel join was not met. A parallel join can only be performed on a pair-wise join of two SPD Server tables, and the query calls three SPD Server tables.

```
LIBNAME path1 sasspds .... IP=YES;

PROC SQL;
create table junk as
select *
  from path1.table1 a,
       path1.table2 b,
       path1.table3 c
 where a.i = b.i and b.i = c.i;
quit;
```

Parallel Join Example 3

Multiple parallel joins can be used in the same SQL query, as long as the SQL planner can perform the query using more than one pairwise join. In the next parallel join example, a more complex query contains a union of two separate joins. Both joins are pair-wise joins of two SPD Server tables. There is a pair-wise join between **table1** and **table2**, and then a pair-wise join between **table3** and **table4** is performed concurrently, using the Parallel Join facility.

```
PROC SQL;
create table junk as
select *
  from path1.table1 a,
       path1.table2 b
 where a.i = b.i
 union

select *
  from path1.table3 c,
       path1.table4 d
 where c.i = d.i;
quit;
```

The required criteria to use the SPD Server Parallel Join facility can be more complex than simply requiring a pair-wise join of two SPD Server tables. The Parallel Join facility can handle multiple character columns, numeric columns, or combinations of character and numeric columns that are joined between pairs of tables. Numeric columns do not need to be of the same width to act as a join key, but character columns must be of the same width in order to be a join key. Columns that are involved in a join cannot be derived from a SAS CASE statement, and cannot be created from character manipulation functions such as SUBSTR,

Parallel Group-By Facility

SPD Server SQL Planner optimizations improve the performance of the more frequent query types used in data mining solutions. One of the SQL planner optimizations integrated into SPD Server is tighter integration of the Parallel Group-By capability. Parallel Group-By is a high performance parallel summarization of data executed using SQL. Parallel Group-By is often used in SQL queries (through the use of sub queries) to apply selection lists for inclusion or exclusion. The tighter integration adds performance benefits to nested Group-By syntax.

Parallel Group-By looks for specific patterns in a query that can be performed using parallel processing summarization. Parallel Group-By works against single tables that are used to aggregate data. Parallel processing summarization is limited to the types of functions it can handle.

The Parallel Group-By support in SPD Server has been expanded in many areas. Parallel Group-By is integrated into the WHERE-clause planner code so that it will boost the capabilities of the SPD Server SQL engine. Any section of code that matches the Parallel Group-By trigger pattern will use it. Some examples of SQL syntax that employ Parallel Group-By technology in SPD Server are:

- [Enhanced Group-By Functions](#)
- [Table Aliases Supported](#)
- [Nested Queries Meet Group-By Syntax Requirements](#)
- [Formatted Parallel Group Select](#)
- [Parallel Group-By SQL Options](#)

Enhanced Group-By Functions: Parallel Group-By now supports the following functions in syntax: COUNT, FREQ, N, USS, CSS, AVG, MEAN, MAX, MIN, NMISS, RANGE, STD, STDERR, SUM, VAR. These functions all can accept the DISTINCT term. The listed functions are the minimum summary functions that are required in order to support the SAS Enterprise Marketing Automation tool suite.

Table Aliases Supported: Table aliases are now supported in SPD Server in order to better support front end tools such as SAS Enterprise Marketing Automation. Tools such as SAS Enterprise Marketing Automation generate SQL queries that use table aliases. Table aliases allow both shorter coding syntax and a method to select a specific column in a query that has two tables that share common column names.

Nested Queries Meet Group-By Syntax Requirements: Since the Parallel Group-By functionality is integrated into the SPD Server WHERE-clause planner, now many sections of queries can take advantage of performance enhancements such as parallel processing. Some common performance enhancements are sub-queries that generate value lists in an IN clause, views that now conform to Parallel Group-By syntax, and views that contain nested Group-By syntax.

General Syntax:

SELECT *'project-list'* **FROM** *'table name'* ;

WHERE [*where_expression*];

GROUP BY [*groupby-list*];

HAVING [*having-expression*];

ORDER BY [*orderby-list*];

project-list

Items must be either column names (which must appear in the groupby-list) or aggregate (summary) functions involving a single column (with the exception of "count(*)" which accepts an asterisk argument. At least one aggregate function must be specified. Project items may be aliased (for example, select avg (salary) *as avg\$al* from ...) and these aliases may appear in any where-expression, having-expression, groupby-list or orderby-list. The following aggregate functions are supported: count, avg, avg distinct, count distinct, css, max, min, nmiss, sum, sum distinct, supportc, range, std, stderr, uss, var. "Mean" is a synonym for "avg". "Freq" and "n" are synonyms for "count" except they do not accept the asterisk argument.

table name

Table names may be one- or two-part identifiers (for example, mytable or foo.mytable), the latter requiring a previous "libref" statement to define the domain identifier (for example, foo).

The *where-expression* is optional.

The optional *groupby-list* must be column names or projected aliases.

The optional *having-expression* must be a boolean expression composed of aggregate functions, groupby columns and/or constants.

The optional *orderby-list* must be projected column names or aliases or numbers which represent the position of a projected item (for example, select a, count (*) order by 2).

Since the Parallel Group-By functionality is integrated into the SPD Server WHERE-clause planner, now many sections of queries can take advantage of performance enhancements such as parallel processing. Some common performance enhancements are sub-queries that generate value lists in an IN clause, views that now conform to Parallel Group-By syntax, and views that contain nested Group-By syntax.

Formatted Parallel Group Select

By default, the columns of a group-by statement are grouped by their unformatted value. SQL pass-through parallel groupby provides the capability to also group data by the columns output data format. For example, you could group by the date column of a table with an input format of mmddyy8 and an output format of monname9. Suppose the column has dates 01/01/04 and 01/02/04. Grouping by the unformatted value would put these dates into two separate groups. However, grouping by the formatted month name, would put these values into the same month grouping of January.

You enable or disable pass-through formatted parallel groupby with the following execute commands:

```
PROC SQL;
connect to sasspds
  (dbq=.....);

/* turn on formatted parallel group-by */
execute(reset fmtgrpssel)
  by sasspds;

select *
from connection
to sasspds
(select dte
  from mytable
  groupby dte);

/* turn off formatted parallel group-by */
execute(reset nofmtgrpssel)
  by sasspds;

select *
from connection
to sasspds
(select dte
  from mytable
  groupby dte);

quit;
```

The example code below is extracted from a larger block of code, whose purpose is to make computations based on user-defined classes of age, such as Child, Adolescent, Adult, and Pensioner. The code uses SQL Parallel Group-By features to create the user-defined classes and then uses them to perform aggregate summaries and calculations.

```
/* Use the parallel group-by feature with the */
/* fmtgrpssel option. This groups the data based */
```

```

/* on the output format specified in the table. */
/* This will be executed in parallel.          */

PROC SQL;
connect to sasspds
  (dbq="&domain"
   serv="&serv"
   host="&host"
   user="anonymous");

/* Explicitly set the fmtgrp sel option */

execute(reset fmtgrp sel)
  by sasspds;

title 'Simple Fmtgrp sel Example';
select *
from connection to sasspds
  (select age, count(*) as count
   from fmtest group by age);

disconnect from sasspds;
quit;

PROC SQL;
connect to sasspds
  (dbq="&domain"
   serv="&serv"
   host="&host"
   user="anonymous");

/* Explicitly set the fmtgrp sel option */

execute(reset fmtgrp sel)
  by sasspds;

title 'Format Both Columns Group Select Example';

select *
from connection to sasspds
  (select
   GENDER format=$GENDER.,
   AGE format=AGEGRP.,
   count(*) as count
   from fmtest
   formatted group by GENDER, AGE);

disconnect from sasspds;

quit;

PROC SQL;
connect to sasspds
  (dbq="&domain"
   serv="&serv"
   host="&host"
   user="anonymous");

/* Explicitly set the fmtgrp sel option */

execute(reset fmtgrp sel)
  by sasspds;

title1 'To use Format on Only One Column With Group Select';
title2 'Override Column Format With a Standard Format';

select *

```

```

from connection to sasspds
  (select
    GENDER format=$1.,
    AGE format=AGEGRP.,
    count(*) as count
  from fmttest
  formatted group by GENDER, AGE);

disconnect from sasspds;

quit;

/* A WHERE-clause that uses a format to subset */
/* data is pushed to the server. If it is not */
/* pushed to the server, the following warning */
/* message will be written to the SAS log: */
/* WARNING: Server is unable to execute the */
/* where clause. */

data temp;
set &domain..fmttest;
  where put
    (AGE,AGEGRP.) = 'Child';
run;

```

The complete code example is found in the [User-Defined Formats section](#) of the SPD Server User's Guide chapter on [SPD Server Formats and Informats](#).

Parallel Group-By SQL Options

SPD Server provides the following Parallel Group-By SQL reset options:

- [GRPSEL/NOGRPSEL](#)
- [FMTGRPSEL/NOFMTGRPSEL](#)
- [SCANGRPSEL/NOSCANGRPSEL](#)

GRPSEL/NOGRPSEL

The **GRPSEL/NOGRPSEL** option enables or disables the SPD Server Parallel Group-By facility.

Usage:

```

/* Disable Parallel Group-By */
execute(reset nogrpSEL)
  by sasspds ;

```

FMTGRPSEL/NOFMTGRPSEL

The **FMTGRPSEL/NOFMTGRPSEL** option enables or disables the SPD Server Parallel Group-By use of formats.

Usage:

```

/* Disable Parallel Group-By */
/* use of formats. */
execute(reset nofMTgrpSEL)
  by sasspds ;

```

SCANGRPSEL/NOSCANGRPSEL

Use the **SCANGRPSEL/NOSCANGRPSEL** option to turn the SPD Server index scan facility on and off. The default SPD Server setting uses the index scan facility.

Usage:

```
/* Disable index scan facility */
execute(reset noscangrp sel)
  by sasspds ;

/* Enable index scan facility */
execute(reset scangrp sel)
  by sasspds ;
```

SPD Server STARJOIN Facility

The SPD Server's enhanced SQL planner includes the STARJOIN facility. The SPD Server STARJOIN facility validates, optimizes, and executes SQL queries on data that is configured in a star schema. Star schemas are composed of two or more normalized dimension tables that surround a centralized fact table. The centralized fact table contains data elements of interest derived from the dimension tables.

In data warehouses with large numbers of tables and millions or billions of rows of data, properly constructed star joins can minimize overhead data redundancy during query evaluation. If the SPD Server STARJOIN facility is not enabled, or if SPD Server SQL does not detect a star schema, then the SQL will be processed using pair-wise joins.

How do star joins differ from pair-wise joins? In SPD Server 4.4, properly configured star joins require only three steps to complete, regardless of the number of dimension tables. SPD Server pair-wise joins require one step for each table to complete the join. If a star schema consisted of 25 dimension tables and one fact table, the STARJOIN is accomplished in three steps; joining the tables in the star schema using pair-wise joins will require 26 steps.

When data is configured in a valid SPD Server star schema, and the STARJOIN facility is not disabled, the SPD Server STARJOIN facility can produce quicker and more processor-efficient SQL query performance than would be realized using SQL pair-wise join queries.

More detailed information is available in the documentation chapter on the [STARJOIN Facility](#).

STARJOIN Options

Use the SPD Server SQL STARJOIN facility options to specify how SPD Server implements Star Joins.

- [NOSTARJOIN](#)
- [STARMAGIC](#)
- [DETAILS](#)

NOSTARJOIN

Use the **NOSTARJOIN** option to disable or enable the SPD Server STARJOIN facility.

Usage

```
execute(reset nostar join=<1/0>)
  by sasspds ;
```

NOSTARJOIN=0 enables the SPD Server STARJOIN facility.

NOSTARJOIN=1 disables the SPD Server STARJOIN facility.

STARMAGIC

Use the **STARMAGIC** option to modify the behavior of the SPD Server STARJOIN and override some internal heuristics in order to favor a particular join strategy in the planner. The vaules are bit flags in the STARJOIN code that can be added together to result in a variety of controls.

Usage

```
execute(reset starmagic=<1/2/4/8/16>)
by sasspds ;
```

STARMAGIC=1 forces all dimension tables to be classified as Phase I tables.

STARMAGIC=2 is currently not used.

STARMAGIC=4 requires an exact match on the FACT composite index in order to meet Phase I conditions for STARJOIN.

STARMAGIC=8 disables the IN-SET STARJOIN strategy. The IN-SET strategy is enabled by default.

STARMAGIC=16 disables the COMPOSITE STARJOIN strategy. The COMPOSITE strategy is enabled by default.

DETAILS

Use the **DETAILS** option to print details about your SPD Server STARJOIN facility settings. All internal STARJOIN debugging information is tied to the "stj\$" DETAILS key. Issuing the "stj\$" reset option displays available information as SPD Server attempts to validate a join sub-tree. The RESET DETAILS="stj\$" option is very useful for debugging STARJOIN and SQL statement execution.

Usage

```
execute(reset details="stj$")
by sasspds ;
```

SPD Server Index Scan

SPD Server SQL provides users with the capability to use lightning-fast index scans on large tables. Rather than scanning entire tables which may have million or billions of rows, SPD Server SQL is able to scan cached index metadata instead of sequentially scanning entire large tables. SPD Server SQL provides enhanced index scan support for the following functions:

min, max, count, nmiss, range uss, css, std, stderr, and var. All of the functions can accept the DISTINCT term as well..

All index scan capabilities listed above are available for both standard SPD Server tables as well as clustered tables, with the exception of the DISTINCT qualifier. The DISTINCT index scan function is not available in clustered tables.

The count(*) function is the only function included with the index scan support enhancement that does *not* require an index on the table. For example,

```
select count(*) from tablename;
```

will return the number of rows in the large table *tablename* without performing a row scan of the table. Table metadata is able to return the correct number of rows. As a result, the response is as fast as an index scan, even on an unindexed table in this case.

Count(*) functions with WHERE-clauses require an index for each column referenced in the WHERE-clause, in order for the index scan feature to provide the performance enhancement. For example, suppose SPD Server table Foo has indexes on numeric columns a and b. The following count(*) functions benefit from SPD Server index scan support:

```

select count(*)
  from Foo
   where a = 1;

select count(*)
  from Foo
   where a LT 4
   and b EQ 5;

select count(*)
  from Foo
   where a in (2,4,5)
   or b in (10,20,30);

```

All functions other than count(*) require an index on function columns in order to exploit the index scan performance savings. Minimal WHERE-clause support is available for these queries, as long as all functions use the same column, and the WHERE-clause is a simple clause that uses the LT, LE, EQ, GE, GT, IN, or BETWEEN operator for that column. For example, suppose that the SPD Server table Bar has indexes on numeric columns x and y. The following SQL submissions will be able to exploit the performance gains of index scans:

```

select min(x),
       max(x),
       count(x),
       nmiss(x),
       range(x),
       count(distinct x)
  from Bar;

select min(x),
       max(x),
       count(x),
       nmiss(x),
       range(x),
       count(distinct x)
  from Bar
   where x between 5 and 10;

select min(x),
       max(x),
       count(x),
       nmiss(x),
       range(x),
       count(distinct x)
  from Bar
   where x gt 100;

select min(x),
       min(y),
       count(x),
       count(y)
  from Bar;

```

If any one function in a statement does not meet the index scan criteria, all functions in that statement will revert to being resolved by table scan instead of index scan. Suppose the SPD Server table Oops has indexes on numeric columns x and y. Column z is not indexed. Then, the SPD Server SQL statement below

```

select min(x),
       min(y),
       count(x),
       count(y),
       count(z)
  from Oops;

```

will be entirely evaluated by table scan; index scanning will not be performed on any of the functions. To take advantage of index scans, the statement above could be resubmitted as

```

select min(x),
       min(y),
       count(x),
       count(y)
from Oops;

select count(y)
from Bar;

```

The functions min(x), min(y), count(x), and count(y) will be evaluated using index scan metadata and will exploit the performance gains. The function count(y) will continue to be evaluated by table scan. The count(*) function can be combined with other functions and benefit from index scan performance gains. Continuing to use the SPD Server table Oops with indexes on numeric columns x and y, the following SPD Server SQL statement will benefit from index scan performance:

```

select min(x),
       range(y),
       count(x),
       count(*)
from Oops;

```

SPD Server Index Scan is an extension to the SPD Server Parallel Group-By Facility. The query must first be accepted by Parallel Group-By to be evaluated for an Index Scan. The section on [SPD Server Parallel Group-by Facility](#) contains more detailed information. When SPD Server utilizes the Index Scan optimization, the following message will be printed to the SAS log:

```

SPDS_NOTE: Metascan used to resolve this query.

```

[Optimizing Correlated Queries](#)

Intelligent storage must have the ability to interpret and process complex requests such as correlated queries. A correlated query is a select expression where a predicate within the query has a relationship to a column that is defined in another scope. Today's business and analytic intelligence tools often generate SQL queries that are nested 3 or 4 layers deep. Queries with cross-nested relationships consume significant processor resources and require more time to complete processing. New algorithms in the SQL Planner of SPD Server implement techniques that significantly improve the performance of correlated queries for patterns that permit query rewrites or query de-correlation.

The SQL Planner improves correlated query performance by changing complex rules about nested relationships into a series of simple steps. SPD Server can process the simple steps much faster than it can process the complex rules that arise with multiple levels of nesting. When a query with multiple levels of nesting is submitted to the SQL Planner, the planner examines the relationships between nested and un-nested sections of the query. When a complex nested relationship is found, the SQL Planner restructures or recodes the SQL query into a simpler form using temporary SPD Server tables.

Development work continues to improve the range of sub-expressions that are addressed by the SPD Server rewrite facility. More information is available in the [SQL Query Rewrite Facility](#) Help section, including SPD Server parameter specifications and SQL RESET options.

[Correlated Query Options](#)

The following are SPD Server SQL options for use with correlated query rewrites:

- [_QRW/NO_QRW](#)
- [_QRWENABLE/NO_QRWENABLE](#)

[_QRW/NO_QRW](#)

Use the [_QRW/NO_QRW](#) option to configure SPD Server to enable or disable the query rewrite facility diagnostic output. Specifying this SPD Server RESET option enables or disables various debugging and tracing outputs from the query rewrite facility. The debugging and tracing outputs are generated when the SPD Server query rewrite facility detects sub-expressions that it rewrites and executes the SQL code. The SQL code produces the intermediate results and final rewritten SQL statement. By default, the SPD Server [_QRW](#) option for diagnostic output is not enabled.

SPD Server provides alternate expressions that do the same thing as the `_QRW/NO_QRW` option. They are the `_QRW=1/_QRW=0` option and the `NO_QRW=0/NO_QRW=1` option.

Usage:

```
/* Enable query rewrite diagnostics */
execute(reset _qrw)
  by sasspds ;

/* A second way to enable      */
/* query rewrite diagnostics */
execute(reset _qrw=1)
  by sasspds ;

/* A third way to enable      */
/* query rewrite diagnostics */
execute(reset no_qrw=0)
  by sasspds ;

/* Disable query rewrite diagnostics */
execute(reset no_qrw)
  by sasspds ;

/* A second way to disable query */
/* rewrite diagnostics          */
execute(reset _qrw=0)
  by sasspds ;

/* Another way to disable query */
/* rewrite diagnostics          */
execute(reset no_qrw=1)
  by sasspds ;
```

[_QRWENABLE/NO_QRWENABLE](#)

Use the `_QRWENABLE/NO_QRWENABLE` option to completely disable the SPD Server query rewrite facility. Disabling the query rewrite facility prevents the rewrite planner from intervening in the SQL flow and from making any optimizing rewrites. This option is not normally specified unless you wish to test if an SQL statement would run faster without rewrite optimization, or if you suspect that the resulting row set that you get from a query rewrite evaluation is incorrect.

SPD Server provides an alternate expression that does the same thing as the `_QRWENABLE/NO_QRWENABLE` option. It is the `_QRWENABLE=1/_QRWENABLE=0` option. The query rewrite facility is enabled in SPD Server by default.

Usage:

```
/* Disable query rewrite */
/* facility              */
execute(reset no_qrwenable)
  by sasspds ;

/* A second way to disable */
/* query rewrite facility */
execute(reset _qrwenable=0)
  by sasspds ;

/* Enable query rewrite */
/* facility              */
execute(reset _qrwenable)
  by sasspds ;
```

```
/* A second way to enable */
/* query rewrite facility */
execute(reset _qrwenable=1)
by sasspds ;
```

Materialized Views

SPD Server allows users to create a SQL view as a materialized view. What makes a materialized view different from an SQL view? For a materialized view, the results of the view statement are computed and saved in a temporary SPD Server table at the time the view is created. For a standard SQL view the results are computed each time the view is referenced in a subsequent SQL statement. As long as there are no changes to any of the input tables that the view is comprised of, the materialized view will return the results from the temporary table when the view is referenced in a SQL statement. If any of the input tables that comprise the view are modified, the materialized view will recompute the results the next time that the view is referenced and it will refresh the temporary table with the new results. The materialized view temporary results table exists for as long as the view is in existence. When a view is dropped or deleted, then the temporary results table is also deleted.

- [Materialized Views Operating Details](#)
- [User Interface for Materialized Views](#)
- [Benefits of Materialized Views](#)
- [Materialized View Example](#)

Materialized Views Operating Details

A materialized view can only be created at the time the SQL view is created. This feature is only available using the SPD Server 4.4 SQL Pass-Through facility. A new keyword Materialized is added to the Create View syntax that identifies the view to be created as a materialized view. When a materialized view is created, the Create View operation will not complete until the temporary results table has been populated. This may add substantial time to the execution of Create View.

Each time a created materialized view is referenced in a SQL statement, there is a check to determine if any of the input tables used to produce the temporary results table have been modified. If not, the temporary table is substituted in place of the view file within the SQL statement. If any of the input tables have been modified, the SQL statement will execute without this substitution so it will act as if it is a standard SQL view reference. There is also a background thread launched at this time that is independent of the SQL statement execution which will refresh the temporary results table. Until this refresh is completed, any incoming references to the view will be treated as standard view references.

Creating a standard SQL view results in a view file being created in the specified domain with the name <viewname>.view.0.0.0.spds9. Creating a materialized view results in an additional SPD Server table being created in the same domain as the view file with the name format <.viewname>.mdfspds9 and corresponding dpf files <.viewname>.dpfspds9. The materialized view table is not visible or accessible to the user by using PROC DATASETS or other SAS procedures. If one or more simple indexes are defined on any of the input tables that are used to create the results table, the indexes will also be created on the materialized view table, as long as the column that was indexed in the input table also exists in the materialized view table.

User Interface for Materialized Views

To create a materialized view, use the following SQL Pass-Through syntax.

```
EXECUTE (Create Materialized View <viewname> as Select ) BY [sasspds | alias];
```

All other references to the view follow the existing SQL syntax, whether it is a standard SQL view or a materialized view. The Materialized keyword is only used in the Create statement. For example, to drop a materialized view, you would use the following syntax.

```
EXECUTE (Drop View <viewname> ) BY [sasspds | alias];
```

If any of the input tables to a materialized view are modified, the next time the view is referenced, a refresh is performed on the materialized view table. You can use an **spdsserv.parm** file option setting to specify the time delay before the materialized view table is refreshed.

```
MVREFRESHTIME=<number-of-seconds> ;
```

Where *<number-of-seconds>* specifies the number of seconds before the refresh will start. You can set the MVREFRESHTIME= option to any integer value between 0 and 86400. The default MVREFRESHTIME= specification is 30 seconds.

The reason that a time delay may be necessary before refreshing a materialized view table is to prevent processor thrashing. Processor thrashing may occur if you refresh the materialized view table when other processes are concurrently processing updates to the tables that are used in the view. If your computing environment does not perform multiple concurrent table updates, then you can set MVREFRESHTIME=0 and eliminate any time delay associated with materialized view refreshes.

Benefits of Materialized Views

Creating a materialized view instead of a Standard SQL view can provide enormous performance benefits when the view is referenced in an SQL statement. For views that contain costly operations such as multiple table joins or operations on very large tables, the execution time for queries containing a materialized view can be orders of magnitude less than a standard view. If the results table produced by the view is relatively small in comparison with the input tables, the execution time for queries using a materialized view may be a few seconds versus several minutes for a standard view.

For example, if it takes on average 20 minutes to produce the result set from a view and the result is in the order of thousands of rows or less, a query referencing the materialized view will now take seconds. Previously using the standard view operations, every time the view was referenced would result in 20 minutes of execution time. The performance benefits should be measured on a case by case basis.

The decision of whether to use a standard view or a materialized view can be primarily driven by how often the input tables to the view are updated versus how often the view is referenced in a SQL statement. If a view is being referenced at least twice before any updates may occur, then the materialized view should provide superior performance. In cases where the defined view can be created very quickly, there is probably not a need for using a materialized view. If the input tables are frequently updated in comparison to how often the view is referenced, a standard view would probably be more efficient.

Materialized View Example

The following code shows the creation and use of a materialized view. The input tables X and Z are created with X having three columns a,b,c and Z having four columns a,b,c,d respectively.

```
data mydomain.X;
  do a = 1 to 1000;
    b = sin(a);
    c = cos(a);
  output;
end;
run;

data mydomain.Z;
  do a = 500 to 1500;
    b = sin(a);
    c = cos(a);
    d = mod(a,99);
  output;
end;
run;

PROC SQL;
connect to sasspds (dbq='mydomain'
  host='myhost'
  serv='myport'
  user='me'
  passwd='mypasswd');

execute (create materialized view XZVIEW as
  select *
    from Z
   where a in
      (select a from X))
  by sasspds;
```

```

select *
  from connection
  to sasspds
  (select *
   from XZVIEW
   where d >90);

execute (drop view XZVIEW);
quit;

```

SPD Server SQL Extensions

SPD Server SQL furnishes several extensions to the SQL language. These extensions are not a part of standardized industry SQL, but they are an integral part of the SPD Server system. These extensions enable systemic data management unique to the SPD Server. The SPD Server SQL uses a special pass-through facility that employs these extensions for data manipulation and extraction operations. The following section discusses the roles of the following extensions which enable SPD Server's SQL pass-through facility. Extensions users should know are [LIBREF statements](#), [LIBREF clauses](#), [BEGIN ASYNC OPERATION statements](#), [END ASYNC OPERATION statements](#), [LOAD statements](#), and [COPY statements](#).

- [BEGIN and END ASYNC OPERATION Statements](#)
- [LOAD Statement](#)
- [COPY Statement](#)

BEGIN and END ASYNC OPERATION Statements

Asynchronous statements are a useful technique you can use to harness the multi-processor power of SPD Server. Asynchronous statements enable execution of multiple, independent threads at the same time. The BEGIN ASYNC OPERATION and END ASYNC OPERATION statements allow you to delimit one or more statements for asynchronous, parallel execution. Since the statements execute in parallel, they must not depend on another, because there is no way to guarantee which statement will finish before another statement executes. SPD Server software initiates thread execution according to the order of the statements in the block.

- [Illegal ASYNC Block Statements](#)
- [Legal ASYNC Block Statements](#)
- [Using LIBREFs in an ASYNC Block Statement](#)
- [Using SQL Options in an ASYNC Block Statement](#)

Usage:

```
execute ([ BEGIN | END ] ASYNCH OPERATION);
```

Illegal ASYNC Block Statements

The statements in this Illegal ASYNC Block example have illegal interdependencies and cannot be expected to work correctly:

```

/* Example of Illegal ASYNC Block Code */

PROC SQL;
  connect to sasspds
  (dbq="my-domain"
   server=host.port
   user='user-name'
   password='user-password'
   other connection options);

```



```

execute(begin async operation)
  by sasspds;

execute(create table T1 as
  select *
  from SRC1)
  by sasspds;

execute(create unique index I1 on
  T1(a,b))
  by sasspds;

execute(end async operation)
  by sasspds;

disconnect from sasspds;
quit;

```

The example violates the interdependency rule. The `create index` statement assumes table `T1` exists and is complete. However, table `T1` is created from table `SRC1`, and may not be complete before the asynchronous create index statement executes. Hence, index `I1` is dependent on a complete table `T1`. The resultant data would not be reliable. The purpose of this example is to illustrate the concept of interdependency, and how **not** to write an ASYNC block.

Legal ASYNC Block Statements

The statements in this Legal ASYNC Block example have no interdependencies.

```

/* Example of Legal ASYNC Block Code          */
/* Creates some tables in the first ASYNC block */
/*                                           */

PROC SQL;
  connect to sasspds
    (dbq="path1"
     server=host.port
     user='anonymous');

  execute(begin async operation)
    by sasspds;

  execute(create table state_al as
    select *
    from allstates
    where state='AL')
    by sasspds;

  execute(create table state_az as
    select *
    from allstates
    where state='AZ')
    by sasspds;
  ...

  execute(create table state_wy as
    select *
    from allstates
    where state='WY')
    by sasspds;

  execute(end async operation)
    by sasspds;

```

```

/*                                     */
/* Create some indexes in the second ASYNC block */
/*                                     */

    execute(begin async operation)
      by sasspds;

    execute(create index county on
             state_al(county))
      by sasspds;

    execute(create index county on
             state_az(county))
      by sasspds;
    ...

    execute(create index county on
             state_wy(county))
      by sasspds;

    execute(end async operation)
      by sasspds;

    disconnect from sasspds;
quit;

```

Why does the second example work correctly? First, each table is created independently. Second, there is a 'synchronization point': the first **END ASYNC** operation. This point ensures that all the tables are created before the second ASYNC statement block begins. (You can also achieve results that are similar to this example by using the [LOAD Statement](#)).

[Using LIBREFs in an ASYNC Block Statement](#)

To refer to a two-part table name inside an ASYNC block, you must re-execute the LIBREF statement issued before entering the block. Conversely, if you issue a LIBREF statement inside the ASYNC block, it does not extend outside the ASYNC block. An ASYNC block creates a distinct scope for the LIBREF. To work correctly, a LIBREF statement must be located inside the ASYNC block, and the LIBREF statement must precede the first SQL statement that references it.

```

/* Example of Legal Code using LIBREFs in an ASYNC Block */
/* Create some tables in the first ASYNC block          */

PROC SQL;
  connect to sasspds
    (dbq="path1"
      server=host.port
      user='anonymous');

  execute(begin async operation)
    by sasspds;

  execute(libref path1 engopt='dbq="path1"
          server=host.port
          user="anonymous"')
    by sasspds;

  execute(libref path2 engopt='dbq="path1"
          server=host.port
          user="anonymous"')
    by sasspds;

  execute(create table path1.southeast as
          select a.customer_id,
                 a.region,

```

```

        b.sales
    from    path1.customer a,
           path2.orders b
    where  a.customer_id = b.customer_id
    and    a.region='SE')
    by sasspds;

    ....

execute(create table path1.northeast as
        select a.customer_id,
               a.region,
               b.sales
    from    path1.customer a,
           path2.orders b
    where  a.customer_id = b.customer_id
    and    a.region='NE')
    by sasspds;

execute(end async operation)
    by sasspds;

disconnect from sasspds;
quit;

```

Using SQL Options in an ASYNC Block Statement

SPD Server SQL options must be set globally for all execute statements in the ASYNC block. These options must be set using an execute statement before the BEGIN ASYNC operation. This example uses code blocks from the preceding example to show how to print a method tree without executing the SQL.

```

/*                                     */
/* Example of Legal SQL Options in ASYNC Block */
/*                                     */

PROC SQL;
    connect to sasspds
        (dbq="path1"
         server=host.port
         user='anonymous');

    execute(reset noexec _method)
        by sasspds;

    execute(begin async operation)
        by sasspds;

    execute(libref path1
        engopt='dbq="path1"
        server=host.port
        user="anonymous"')
        by sasspds;

    execute(libref path2
        engopt='dbq="path1"
        server=host.port
        user="anonymous"')
        by sasspds;

    execute(create table path1.southeast as
        select a.customer_id,
               a.region,
               b.sales
    from    path1.customer a,
           path2.orders b

```

```

        where a.customer_id = b.customer_id
        and   a.region='SE')
        by sasspds;

        ....

execute(create table path1.northeast as
        select a.customer_id,
               a.region,
               b.sales
        from   path1.customer a,
               path2.orders b
        where  a.customer_id = b.customer_id
        and   a.region='NE')
        by sasspds;

execute(end async operation)
        by sasspds;

disconnect from sasspds;
quit;

```

LOAD Statement

The LOAD statement enables table creation (with one or more indexes) with a single statement. The data source for the statement is a SELECT clause. The **SELECT** list in the clause defines the columns for the new table. All characteristics of the columns (variables) in the select list are preserved, becoming permanent attributes of the new table's column definitions. The target table for the LOAD TABLE statement must be on the local machine.

Usage:

```

execute (LOAD TABLE table spec
        < WITH index spec
        < WITH index spec >>
        by sasspds;

```

In the following example, each execute statement creates a table for one U.S. state using a global table called STATE that contains many states. The first execute statement uses LOAD to create table STATE_AL (Alabama), and creates an index on the COUNTY column. The structure of the state table STATE_AL and the data in the state table both come from the global table STATE. The data in STATE_AL is the subset of all records from STATE where the STATE column variable equals 'AL'. LOAD creates a table for states from Alabama to Wyoming, with each state's table indexed by county and mirroring the structure of the parent table STATE.

```

execute(load table state_al
        with index county
        on (county) as
        select *
        from state
        where state='AL')
        by sasspds;

execute(load table state_az
        with index county
        on (county) as
        select *
        from state
        where state='AZ')
        by sasspds;

...

execute(load table state_wy
        with index county

```

```

on (county) as
select *
from state
where state='WY')
by sasspds;

```

In general, the LOAD statement is faster than a corresponding "create table" / "create index" statement pair, because it builds the table and associated index(es) asynchronously using parallel processing.

COPY Statement

The COPY table statement creates a copy of a SPD Server table with or without the table index(es). For the COPY table statement to work, the source and target tables must be on the local machine. By default, the software creates an index(es). The COPY table statement is faster than either of the following CREATE and LOAD statements:

```

create table ...
as select ...
create index ...

```

or

```

load table ...
with index...
as select ...

```

The COPY statement is faster than the two above statements because it uses a more direct access path than the SQL SELECT clause when accessing the data.

In the example that follows, two new tables are created: T_NEW and T2_NEW. The first table, T_NEW, is created with index structures identical to table T_OLD. The second table, T2_NEW, is unindexed regardless of the structure of table T2_OLD.

```

execute(copy table t_new
        from t_old)
by sasspds;

execute(copy table t2_new
        from t2_old
        without indexes)
by sasspds;

```

The COPY statement also supports an ORDER BY clause you use to create a new table with a sort order on one or more columns of the new table. While COPY TABLE does not support all of the options of PROC SORT, you can achieve substantial performance gains when creating this ordered table by using COPY with an ORDER BY clause when appropriate.

The next example copies the table T_OLD to T_NEW using the order by clause. The data will be ordered by columns: x in ascending order, y in descending order, and z in ascending order. The results are the same if you run PROC SORT on the columns using the same BY clause. The syntax of the COPY ORDER BY follows the normal SQL ORDER BY clause, but the column identifiers that you can specify are restricted. You can only specify actual table columns when using the COPY ORDER BY clause.

```

execute(copy table t_new
        from t_old
        order by x, y desc, z asc)
by sasspds;

```

Differences between SAS SQL and SPD Server SQL

This section overviews some of the functional differences between SAS SQL and SPD Server SQL. A great deal of SAS SQL functionality is integrated into SPD Server. Exceptions between SAS and SPD Server SQL are listed below.

- [Reserved Keywords](#)
- [Table Options and Delimiters](#)
- [Mixing Scalar Expressions and Boolean Predicates](#)
- [INTO Clause](#)
- [Tilde Negation](#)
- [Nested Queries](#)
- ["USER" Value](#)
- [Supported Functions](#)

Reserved Keywords

SPD Server uses keywords to initiate statements or refer to syntax elements. For example, the words "where" and "group" can only be used in certain ways because there are WHERE and GROUP BY clauses. Keywords are treated as reserved words. That means you cannot use keywords when naming a LIBREF, a table, a column or an index.

In contrast, SAS allows keywords in some, but not all, syntax locations. The documentation chapter [SPD Server SQL Syntax Reference Guide](#) contains a list of [reserved SPD Server SQL keywords](#).

Table Options and Delimiters

SPD Server SQL uses brackets to delimit table options. SAS SQL uses parentheses as delimiters. You can place table options in a "create table" statement. You must put table options within parentheses to delimit column definitions within a table.

Mixing Scalar Expressions and Boolean Predicates

SPD Server SQL does not allow mixing scalar expressions with Boolean predicates. SAS SQL does allow mixing scalar expressions with Boolean predicates in most places. The Help section on [Scalar Expressions and Boolean Predicates](#) contains more information on permissible expression content.

INTO Clause

SPD Server SQL does *not* support the INTO clause, as in

```
select a, b into :var1, :var2 from t where a > 7;
```

In contrast, SAS SQL supports the INTO clause.

Tilde Negation

SPD Server SQL supports the use of the tilde only to negate the 'equals' operator, "~=" (not equals). SAS SQL allows broader use of the tilde (~) character, where the tilde is synonymous with "not" and can be combined with various operators. For example, SAS SQL can use the tilde with 'between' "~ between" (not between). SPD Server does not recognize that expression.

Nested Queries

SAS SQL permits sub-queries without parentheses delimiters in more places than SPD Server SQL. SPD Server SQL uses parentheses to explicitly group sub-queries or expressions that are nested within a query statement whenever possible. Queries with nested expressions execute more reliably and are also easier to read.

"USER" Value

SAS SQL permits sub-queries without parentheses delimiters in more places than SPD Server SQL. SPD Server SQL uses parentheses whenever possible to explicitly group sub-queries or expressions that are nested within a query statement. Queries with nested expressions execute more reliably and are easier to read.

SPD Server SQL does not support the "USER" keyword in the INSERT statement. For example, the following query will fail in SPD Server SQL:

```
insert into t1(myname) values(USER);
```

Supported Functions

SPD Server SQL supports the following functions:

abs, addr, arcos, arsin, atan, band, betainv, blshift, bnot, bor, brshift, bxor, byte, ceil, cinv, collate, compbl, compound, compress, cos, cosh, css, cv, daccdb, daccdbsl, daccsl, daccsyd, dacctab, date, datejul, datepart, datetime, day, dcss, depdb, depdbsl, depsl, depsyd, deptab, dequote, dhms, digamma, dmax, dmean, dmin, drange, dstd, dstderr, dsum, duss, dvar, erf, erfc, exp, finv, fipname, fipname1, fipstate, floor, fnonmiss, fuzz, gaminv, gamma, hms, hour, int, intck, intnx, intrr, irr, ispexec, isplink, kurtosis, left, length, lgamma, log, log10, log2, lowercase, max, mdy, mean, min, minute, mod, month, mort, n, netpv, nmiss, npv, ordinal, poisson, probbeta, probbnml, probchi, probf, probgam, probhypr, probit, probnegb, probnorm, probt, qtr, quote, range, ranuni, rank, recip, repeat, reverse, right, round, saving, second, sign, signum, sin, sinh, skewness, sqrt, std, stderr, stfips, stname, stname1, substr, sum, tan, tanh, time, timepart, tinv, today, tranwrdr, trigamma, trim, upcase, uss, var, weekday, year, zipfips, zipname, zipname1, and zipstate.

Ranuni functions may show slight variation from run to run due to the impact of parallel processing.

Note that **date, int, left, right** and **trim** are reserved keywords; therefore, they must be preceded with a backslash in SPD Server SQL queries:

```
select \date() from t ;
```

Copyright © 2008 by SAS Institute Inc., Cary, NC, USA. All rights reserved.

SAS Scalable Performance Data Server SQL Syntax Reference Guide

Contents

- [Overview](#)
- [Document Conventions](#)
- [SQL Syntax Definitions](#)
- [SQL Statements](#)
 - [Alter Table Statement](#)
 - [Connect Statement](#)
 - [Create Index Statement](#)
 - [Create Table Statement](#)
 - [Create View Statement](#)
 - [Delete Statement](#)
 - [Describe Table Statement](#)
 - [Describe View Statement](#)
 - [Disconnect Statement](#)
 - [Drop Index Statement](#)
 - [Drop Table Statement](#)
 - [Drop View Statement](#)
 - [Execute Statement](#)
 - [Insert Statement](#)
 - [Reset Statement](#)
 - [Select Statement](#)
 - [Update Statement](#)
 - [Validate Statement](#)
- [NEW SQL Statements](#)
 - [Async Operation Statement](#)
 - [Contents Statement](#)
 - [Copy Table Statement](#)
 - [Create Materialized View Statement](#)
 - [LIBREF Statement](#)
 - [Load Table Statement](#)
- [SQL Building Blocks](#)
 - [Alias Name](#)
 - [Atomic Expression](#)
 - [Between Predicate](#)
 - [Boolean Expression](#)
 - [Case Expression](#)
 - [Column Definition](#)
 - [Column Definition List](#)
 - [Column Modifier](#)
 - [Column Name](#)
 - [Column Name List](#)
 - [Column Specifications](#)
 - [Comparative Operators](#)
 - [Comparison Predicates](#)
 - [Connection String](#)
 - [Constant](#)
 - [Contains Predicate](#)
 - [Data Types](#)
 - [Date / Time String](#)
 - [DBMS Options](#)
 - [Digits \(Numeric\)](#)

- [Exists Predicate](#)
- [Function Arguments](#)
- [Function Expressions](#)
- [Function Name](#)
- [Identifier](#)
- [In Predicate](#)
- [Index Name](#)
- [Insert Set List](#)
- [Insert Source](#)
- [Insert Value](#)
- [Insert Values List](#)
- [Letter \(Alpha\)](#)
- [LIBREF Name](#)
- [LIKE Predicate](#)
- [Missing Value](#)
- [Null Predicate](#)
- [Number](#)
- [ORDER BY Clause](#)
- [Pass-Through Spec](#)
- [Predicate Types](#)
- [Quantified Comparison Predicate](#)
- [Query Expression](#)
- [Query Spec](#)
- [Scalar Expression](#)
- [Select Item](#)
- [Select Spec](#)
- [Set Value List](#)
- [Soundslike Predicate](#)
- [String](#)
- [Subquery](#)
- [Table Alias](#)
- [Table Join](#)
- [Table Name](#)
- [Table Options](#)
- [Table Reference](#)
- [Table Spec](#)
- [Truth Value](#)
- [Unsigned](#)
- [WHERE Clause](#)
- [With Index Spec](#)

[Overview](#)

This chapter describes the SQL syntax that is allowed with the Scalable Performance Data (SPD) Server. SPD Server SQL is a dialect of SQL. That is, it combines SQL-92, SAS SQL and extensions that are specific to SPD Server. Whenever possible, SPD Server attempts to conform to SAS SQL.

[Document Conventions](#)

- [Productions](#)
- [Production Links / References](#)
- [Literal Text](#)
- [Optional Text](#)

- [Selection Lists](#)

[Productions](#) - The syntax uses building blocks which are referred to as "productions". Productions are denoted by the symbol " ::= ". To the left of the equal sign is a production name; to the right of the equal, or on the next line, is a list of production constructs. If a production has more than one possible construct, the alternatives are separated by a vertical bar "|". Read productions top-down. For example, reading the delete statement, there are literal keywords and two subproductions, a "table_spec" and the "WHERE clause".

[Production Links / References](#) - Subproductions that are referenced within a production definition are HTML links to their definitions. You can navigate the links with an HTML browser.

[Literal Text](#) - Traversing down a syntax tree leads to leaf/terminal definitions. The definitions are composed either of keywords (select), identifiers (names of tables, columns, etc.) or symbols (punctuation, operators, etc.). Keywords and identifiers are shown with bold, capitalized text. In contrast, symbols are shown with single quotation marks and are bold.

[Optional Text](#) - Optional syntax is delimited by square brackets, "[" and "]". Optional lists (syntax elements that are repeated) are denoted by "[" and "]"*". The "*" signifies zero or more occurrences of the bracketed syntax.

[Selection Lists](#) - Selection lists, that allow you to choose from a list of alternative syntax elements, are denoted by braces "{" and "}". These elements are separated by a vertical bar "|". The selection list itself is **not** optional; you must choose at least one element. If you must choose one or more of the elements, the list is terminated with a "}" + ". The "+" indicates one or more occurrences of the delimited syntax.

Note: The browser displays links best with underscores. To view underscores using Netscape, refer to the option under the File Command: Options/General Preferences/Appearance.

[SQL Syntax Definitions](#)

- [Statement \(Query\)](#)
- [Scalar Expressions Contrasted with Boolean Predicates](#)
- [Strings](#)
- [Identifiers](#)
- [\(Reserved\) Keywords](#)

[Statement \(Query\)](#)

One or more syntax elements terminated by a semicolon.

[Scalar Expressions Contrasted with Boolean Predicates](#)

Scalar expressions represent a single data value, either a numeric or a string from a constant specification. Examples include: 1, 'hello there', '31-DEC-60'd), a function (that is, "avg(a*b)"), a column/variable (that is foo.bar), the case expression, or even a subquery which returns a single run-time value. Boolean predicates are either "true" or "false". They are used in WHERE clauses, having clauses and in the case expression. You cannot select predicates, nor can you assign them to columns (that is, in an update statement). Scalar expressions and Boolean predicates **cannot** be used interchangeably, although SAS SQL does allow you to mix the expressions.

[Strings](#)

SPD Server SQL strings are character streams which are delimited by either single or double quotation marks. If you use a single quotation mark to begin a string, you must use a single quotation mark to terminate the string. To embed a single quotation mark in a string, use two single quotation marks together. For example,

```
SELECT 'it's a wonderful life' from mytable.
```

You can use double quotation marks in the same manner. There is another way to embed a single quotation mark without doubling the character. You can use double-quotation marks as delimiters. For example,

```
SELECT "it's a wonderful life" from mytable.
```

In some of the syntax specifications that follow, a "user-defined" or "database-specific" string is noted. Delimit these strings with a bracket or parenthesis. Characters between the delimiters are considered part of the string up to, but not including, the matching delimiter.

```
CONNECT to sasspds(
  user='john'
  passwd='foobar'
  options=(a b c)
);
```

The dbms_options string is

```
user='john'
passwd='foobar'
options=(a b c).
```

In this example, the first right-parenthesis is considered part of the string. It is not the *matching* termination delimiter.

Identifiers

Identifiers are the names of librefs, tables, indexes and columns as well as table and column aliases.

(Reserved) Keywords

Keywords are used to initiate statements and syntax elements. For example, WHERE or GROUP BY clauses. Keywords are also "reserved". They cannot be used for identifiers because this use introduces ambiguity. For example, "select unique unique unique from from from;" is a valid but ambiguous statement. Below is a list of current SPD Server keywords. Some words have been reserved for future enhancements to SPD Server SQL:

add, all, alter, and, any, as, asc, async, begin, between, both, by, calculated, cascade, case, char, character, column, connect, connection, contains, contents, copy, corr, corresponding, create, cross, date, dec, decimal, default, delete, desc, describe, dictionary, disconnect, distinct, double, drop, else, end, engname, engopt, eq, except, execute, exists, false, float, for, format, from, full, ge, grant, group, gt, having, in, index, indexes, informat, inner, insert, int, integer, intersect, into, is, join, label, le, leading, left, libref, like, load, lower, lt, match, missing, modify, natural, ne, no, not, notin, null, num, numeric, on, operation, option, or, order, outer, overlaps, partial, precision, privileges, public, real, references, reset, restrict, revoke, right, select, set, smallint, some, table, then, to, trailing, trim, true, union, unique, unknown, update, upper, using, validate, values, varchar, verbose, view, when, where, with, without, yes

SQL Statements

- [Alter Table Statement](#)
- [Connect Statement](#)
- [Create Index Statement](#)
- [Create Table Statement](#)
- [Create View Statement](#) - Create a view upon one or more tables
- [Delete Statement](#) - Delete records
- [Describe Table Statement](#) - Describe a table definition
- [Describe View Statement](#) - Describe a view definition
- [Disconnect Statement](#) Pass-Through Statement
- [Drop Index Statement](#) - Drop an index from a table
- [Drop Table Statement](#) - Drop a table definition
- [Drop View Statement](#) - Drop a view definition
- [Execute Statement](#) - Pass-Through Statement
- [Insert Statement](#) - Add records
- [Reset Statement](#) - Reset session options and flags
- [Select Statement](#) - Retrieve information
- [Update Statement](#) - Update records
- [Validate Statement](#) - Validate a given select specification

Alter Table Statement

The Alter table statement changes a table definition.

```
alter table statement ::=
  ALTER TABLE table spec
  { { ADD|MODIFY|ALTER [ COLUMN ] column def list } |
    { DROP [ COLUMN ] column name list }
  }+ ';'

```

Connect Statement

The Connect statement creates a pass-through connection.

```
connect statement ::=
  CONNECT TO libref name [ [ AS ]
  alias name ] '('
  dbms options ')' ] ';'

```

Create Index Statement

The Create Index statement creates an index on a table.

```
create index statement ::=
  CREATE [ UNIQUE ] INDEX index name ON
  table spec '(' column name list ')' ';'

```

Create Table Statement

The Create Table statement creates a table definition.

```
create table statement ::=
  CREATE TABLE table spec
  { '(' column def list ')' | AS
  select spec | LIKE
  table spec } ';'

```

Create View Statement

```
create view statement ::= CREATE VIEW
table spec AS
select spec ';'

```

Delete Statement

```
delete statement ::= DELETE FROM
table spec [
where clause ] ';'

```

Describe Table Statement

```
describe table statement ::=
  DESCRIBE TABLE table spec [ [',' ]
    table spec ]* ';'

```

Describe View Statement

```
describe view statement ::=
  DESCRIBE VIEW table spec [ [',' ]
    table spec ]* ';'

```

Disconnect Statement

```
disconnect statement ::= DISCONNECT FROM
libref name ';'

```

Drop Index Statement

```
drop index statement ::=
  DROP INDEX index name [ [',' ]
    index name ]* FROM
    table spec ';'

```

Drop Table Statement

```
drop table statement ::= DROP TABLE
table spec [ [',' ]
table spec ]* ';'

```

Drop View Statement

```
drop view statement ::=
  DROP VIEW table spec [ [',' ]
    table spec ]* ';'

```

Execute Statement

```
execute statement ::= EXECUTE '('
passthru spec ')' BY
libref name ';'

```

Insert Statement

```
insert statement ::=
    INSERT INTO table spec [ '('
        column name list ')' ]
    insert source ';'

```

Reset Statement

```
set option statement ::=
    { SET OPTION | RESET }
    { identifier
    [ '=' { constant |
identifier |
truth value
    | DEFAULT } ] }+

```

Select Statement

```
select statement ::=
select spec ';'

```

Update Statement

```
update statement ::=
    UPDATE table spec
    SET column name '='
scalar expr [ ','
column name '='
scalar expr ]*
    [ where clause ] ';'

```

Validate Statement

```
validate statement ::= VALIDATE
select spec ';'

```

NEW SQL Statements

- [Async Operation Statement](#) — Delimit an asynchronous execution block
- [Contents Statement](#) — Perform a SAS "proc contents" on a table
- [Copy Table Statement](#) — Copy a table, and optionally no indexes on the table, to another table on the same local machine
- [Create Materialized View Statement](#) — create a SQL view as a materialized view. In a materialized view, the results of the view statement are computed and saved in a temporary SPD Server table when the view is created.
- [LIBREF Statement](#) — Perform a SAS LIBREF assignment
- [Load Table Statement](#) — Create a table, and optionally indexes on the table on the local machine, with a select statement

New SQL Statements

Async Operation Statement

```
async operation statements ::= { BEGIN | END } ASYNC OPERATION ';' 
```

Contents Statement

```
contents statement ::= CONTENTS  
table spec [ VERBOSE ] ';' 
```

Copy Table Statement

```
copy table statement ::=  
COPY TABLE table spec FROM  
table spec [ WITHOUT INDEXES ] [ORDER BY  
column name  
[ ASC | DESC ] [ ','  
column name [ ASC | DESC ]]] ';' 
```

Create Materialized View Statement

```
create materialized view statement ::= CREATE MATERIALIZED VIEW  
  
table spec AS  
  
select spec ';' 
```

LIBREF Statement

```
libref statement ::=  
LIBREF libref name [ ENGNAME '='  
identifier ] [ ENGOPT '='  
string ] ';' 
```

Load Table Statement

```
load table statement ::=  
LOAD TABLE table spec [ WITH  
with index spec [ ','  
with index spec ]* ]  
AS select spec ';' 
```

SQL Building Blocks

- [Alias Name](#)
- [Atomic Expression](#)
- [Between Predicate](#)
- [Boolean Expression](#)
- [Case Expression](#)
- [Column Definition](#)
- [Column Definition List](#)
- [Column Modifier](#)
- [Column Name](#)
- [Column Name List](#)
- [Column Specifications](#)
- [Comparative Operators](#)
- [Comparison Predicates](#)
- [Connection String](#)
- [Constant](#)
- [Contains Predicate](#)
- [Data Types](#)
- [Date / Time String](#)
- [DBMS Options](#)
- [Digits \(Numeric\)](#)
- [Exists Predicate](#)
- [Function Arguments](#)
- [Function Expressions](#)
- [Function Name](#)
- [Identifier](#)
- [In Predicate](#)
- [Index Name](#)
- [Insert Set List](#)
- [Insert Source](#)
- [Insert Value](#)
- [Insert Values List](#)
- [Letter \(Alpha\)](#)
- [LIBREF Name](#)
- [LIKE Predicate](#)
- [Missing Value](#)
- [Null Predicate](#)
- [Number](#)
- [ORDER BY Clause](#)
- [Pass-Through Spec](#)
- [Predicate Types](#)
- [Quantified Comparison Predicate](#)
- [Query Expression](#)
- [Query Spec](#)
- [Scalar Expression](#)
- [Select Item](#)
- [Select Spec](#)
- [Set Value List](#)
- [Soundslike Predicate](#)
- [String](#)
- [Subquery](#)
- [Table Alias](#)
- [Table Join](#)
- [Table Name](#)
- [Table Options](#)
- [Table Reference](#)
- [Table Spec](#)

- [Truth Value](#)
 - [Unsigned](#)
 - [WHERE Clause](#)
 - [With Index Spec](#)
-

SQL Building Blocks

Alias Name

```
alias name ::=
  identifier
```

Atomic Expression

```
atomic expr ::=
  constant |
  column spec
```

Between Predicate

```
between pred ::=
  scalar expr [ NOT ] BETWEEN
  scalar expr AND
  scalar expr
```

Boolean Expression

```
boolean expr ::=
  | [ NOT ] { predicate | '('
  boolean expr ')' } [ IS [ NOT ]
  truth value ]
  | boolean expr { AND | OR }
  boolean expr
```

Case Expression

```
case expr ::=
  CASE { WHEN boolean expr THEN
  scalar expr }+ [ ELSE
  scalar expr ] END
  | CASE scalar expr { WHEN
  scalar expr THEN
  scalar expr }+ [ ELSE
  scalar expr ] END
```

Column Definition

```
column def ::=  
  column name  
  data type [  
  column modifier ]* [ NOT NULL ]
```

Column Definition List

```
column def list ::=  
  column def [ ','  
  column def ]*
```

Column Modifier

```
column modifier ::=  
  FORMAT '=' <quoted or nonquoted SAS format specification>  
  | LABEL '=' string
```

Column Name

```
column name ::=  
  identifier
```

Column Name List

```
column name list ::=  
  column name [ [',']  
  column name ]*
```

Column Specifications

```
column spec ::=  
  [ CALCULATED ] column name  
  | table alias '.'  
  column name
```

Comparative Operators

```
comp operator ::=  
  | EQ | '='  
  | NE | '^=' | '~=' | '!=' | '<>'  
  | LT | '<'  
  | GT | '>'  
  | LE | '<='  
  | GE | '>='
```

Comparison Predicates

```
comparison pred ::=
  scalar expr {
  comp operator
  scalar expr }+
```

Connection String

```
connection string ::= <user-defined
string delimited by ending/matching parenthesis>
```

Constant

```
constant ::=
  | number | missing value
  | string | date/time string
  | NULL
```

Contains Predicate

```
contains pred ::=
  scalar expr { CONTAINS | '?' }
  scalar expr
```

Data Types

```
data type ::=
  { CHAR[ACTER] | VARCHAR } [ '('unsigned ')' ]
  | { INT[EGER] | SMALLINT }
  | { NUM[ERIC] | DEC[IMAL] | FLOAT }
  [ '(' unsigned [ ',' unsigned ] ')' ]
  | REAL | DOUBLE PRECISION | DATE
```

Date / Time String

```
date/time string ::=
  string{D|T|DT}
```

DBMS Options

```
dbms options ::= <user-defined
string delimited by ending/matching parenthesis>
```

Digits (Numeric)

```
digit ::= '0' <through> '9'
```

Exists Predicate

```
exists pred ::= EXISTS subquery
```

Function Arguments

```
function args ::=  
  scalar expr [ ',' scalar expr ]* | DISTINCT scalar expr | [ DISTINCT ] '*'
```

Function Expressions

```
function expr ::=  
func name '('  
function args ')'
```

Function Name

```
function name ::=  
identifier
```

Identifier

```
identifier ::= ['\']{  
letter|<underscore>}{  
letter|  
digit|<underscore>}*
```

In Predicate

```
in pred ::=  
  scalar expr { [ NOT ] IN | NOTIN } {  
  subquery | '('  
  constant [ ','  
  constant ]* ')'
```

Index Name

```
index name ::=
```

[identifier](#)

Insert Set List

```
insert set list ::= SET
set value list [ SET
set value list ]*
```

Insert Source

```
insert source ::=
| insert values list
| insert set list
| query expr
```

Insert Value

```
insert value ::= VALUES '('
scalar expr [ ','
scalar expr ]* ')'
```

[Insert Values List](#)

```
insert values list ::=
insert value [
insert value ]*
```

Letter (Alpha)

```
letter ::= 'a' <through> 'z' <or> 'A' <through> 'Z'
```

LIBREF Name

```
libref name ::=
identifier
```

LIKE Predicate

```
like pred ::=
scalar expr [ NOT ] LIKE
scalar expr
```

Missing Value

```
missing value ::= '.'[  
letter]
```

Null Predicate

```
null pred ::=  
scalar expr IS [ NOT ] { NULL | MISSING }
```

Number

```
number ::=  
  {unsigned | {  
    digit}+'.'[{  
    digit}+]|'.'{  
    digit}+}[{'e'|'E'}[ '+' | '-' ]{  
    digit}+]
```

ORDER BY Clause

```
order by clause ::=  
  ORDER BY atomic expr [ ASC | DESC ] [ ','  
  atomic expr [ ASC | DESC ] ]*
```

Pass-Through Spec

```
passthru spec ::=  
  <database-specific string delimited by ending/matching parenthesis>
```

Predicate Types

```
predicate ::=  
  | comparison pred  
  | between pred  
  | in pred  
  | like pred  
  | null pred  
  | quantified comparison pred  
  | exists pred  
  | contains pred  
  | soundslike pred
```

Quantified Comparison Predicate

```
quantified comparison pred ::=
  scalar expr
  comp operator { ALL | SOME | ANY }
  subquery
```

Query Expression

```
query expr ::=
  query spec
  | query expr { [ OUTER ] UNION | EXCEPT | INTERSECT } [ CORRESPONDING ] [ ALL ]
  query expr
```

Query Spec

```
query spec ::=
  SELECT [ DISTINCT | UNIQUE ] select item [ ','
  select item ]*
  FROM table ref [ ','
  table ref ]*
  [ WHERE boolean expr ]
  [ GROUP BY scalar expr [ ','
  scalar expr ]* ]
  [ HAVING boolean expr ]
```

Scalar Expression

```
scalar expr ::=
  | atomic expr
  | function expr
  | '(' scalar expr ')'
```

```
scalar expr ::=
  | subquery
  | scalar expr { '+' | '-' | '*' | '/' | '|' | '**' }
scalar expr
  | { '+' | '-' } scalar expr
  | case expr
```

Select Item

```
select item ::=
  '*'
  | identifier '.' '*'
  | scalar expr [ [ AS ]
  identifier ] [
  column modifier ]*
```

Select Spec

```
select spec ::=
query expr [
order by clause ]
```

Set Value List

```
set value list ::=
column name '='
scalar expr [ ','
column name '='
scalar expr ]*
```

Soundslike Predicate

```
soundslike pred ::=
scalar expr '='*
scalar expr
```

String

```
string ::=
<a single- or double-quoted
literal string -- see Strings>
```

Subquery

```
subquery ::= '('
query expr ')'
```

Table Alias

```
table alias ::=
identifier
```

Table Join

```
table join ::=
table ref [ INNER | { LEFT | RIGHT | FULL }
[ OUTER ] ] JOIN table ref
{ ON boolean expr | USING '('
column name list ')'}
| '(' table join ')'
```

Table Name


```
table name ::=  
identifier
```

Table Options

```
table options ::= <user-defined  
string delimited by ending/matching bracket>
```

Table Reference

```
table ref ::=  
  table spec [ [ AS ]  
  identifier ]  
  | subquery [ [ AS ]  
  identifier ] [ '('  
  column name list ')' ]  
  | CONNECTION TO identifier '('  
  connection string ')' [ [ AS ]  
  identifier ]  
  | table join
```

Table Spec

```
table spec ::=  
  | table name [ '['  
  table options ']' ]  
  | libref name '.'  
  table name [ '['  
  table options ']' ]
```

[Truth Value](#)

```
truth value ::= { TRUE | YES } | { FALSE | NO }
```

Unsigned

```
unsigned ::= {  
digit }+
```

WHERE Clause

```
where clause ::= WHERE  
boolean expr
```

With Index Spec

```
with index spec ::= [ UNIQUE ] INDEX  
index name ON '('  
column name list ')'
```

[Copyright © 2008 by SAS Institute Inc., Cary, NC, USA. All rights reserved.](#)

SAS Scalable Performance Data (SPD) Server SQL Access Library API Reference

- [Introduction](#)
 - [Overview of SPQL Usage](#)
 - [SPQL API Description](#)
 - [SPQL API Functions](#)
 - [SPQL Function Return Codes](#)
-

[Introduction](#)

This chapter describes the Scalable Performance Data Server SQL access library API (Application Programming Interface) and provides some simple examples. This chapter refers to the Scalable Performance Data Server SQL access library as SPQL. Read this chapter if you want a library that provides a C-language compatible interface to write user applications to access an SPD Server SQL server. Because the library was designed for multi-threaded applications, the code is thread safe.

[Overview of SPQL Usage](#)

SPQL enables you to write application: programs that can connect to and access Scalable Performance Data Server (SPD Server) hosts using the SQL language. SPQL is oriented toward connection, allowing you to submit SQL statements to one or more SPD Server SQL servers which execute SQL statements on your behalf.

[SPQL API Description](#)

The C-language H file *spql.h* is provided for customer-written applications. This chapter describes the API functions, their use, and restrictions.

[SPQL API Functions](#)

The SPQL API functions include

- [spqlinit\(\)](#)
- [spqlterm\(\)](#)
- [spqlconnect\(\)](#)
- [spqldisconnect\(\)](#)
- [spqlperform\(\)](#)
- [spqlfreestok\(\)](#)
- [spqltabinfo\(\)](#)

- [spqcolinfo\(\)](#)
- [spqlfetch\(\)](#)
- [spqlgmsg\(\)](#)

[spqlinit\(\)](#)

Initializes the SPQL library for operation.

```
int spqlinit(void)
```

Usage: Performs a one-time initialization which enables the SPQL library to function. For this reason, you must call `spqlinit()` at least once to activate an SPQL program. Do not make other SPQL API calls before calling this function. If you do, the results are unpredictable. When `spqlinit()` successfully completes, you can safely proceed to use the SPQL API in a multi-threaded context.

Note: `Spqlinit()`-is not a thread-safe function. Call it only within a single-threaded context in your application. Alternatively, call it within an application-controlled mutex region.

Parameters: None

Returns: 0 if successful; `SPQL_INITFAILED` if the initialization fails.

[spqlterm\(\)](#)

Is the termination counterpart of the `spqlinit()` function.

```
int spqlterm(void)
```

Usage: Terminates the SPQL library session, disconnecting all active SPD Server SQL server connections and freeing up the memory resources associated with the SPQL run-time library executables.

Parameters: None

Returns: 0 if successful.

[spqlconnect\(\)](#)

Establishes a connection to a specified SPD Server SQL server.

```
int spqlconnect(char *constr, void **contok)
```

Usage: Establishes a connection to the SPD Server SQL server. The **constr** parameter specifies all the connection information needed to establish the connection to an SPD Server SQL server. When a connection is made successfully, a connection, token (**contok**) is returned to the caller.

Parameters:

char *constr

A null-terminated string identifying the SPD Server SQL server to connect to for this session. The syntax for the string is identical to that used for the *SAS PROC SQL* pass-through CONNECT statement.

void **contok

Returns a connection token if the connection successfully completes. You must retain the token; use it in subsequent SPQL library operations that you perform using the connection.

Returns: 0 if successful; SPQL_NOMEM if unable to allocate memory for the connection token; SPQL_CONFAILED if unable to connect successfully to the SPD Server SQL server.

[spqldisconnect\(\)](#)

Terminates a connection from the SPD Server SQL server specified with an `spqlconnect()`.

```
int spqldisconnect(void *contok)
```

Usage: Disconnects from a specified SPD Server SQL server. The caller passes the connection token which was returned from an `spqlconnect()` call. Then, the SPD Server SQL server associated with the connection is disconnected from the caller, and the memory associated with connection token is returned to the system.

Parameters:

void *contok

Connection token previously obtained from `spqlconnect()`.

Returns: 0 if successful.

[spqlperform\(\)](#)

Submits an SQL statement for execution on a given connection.

```
int spqlperform(void *contok, char *stmtbuf, int stmtlen,  
               int *actions, void **stmttok);
```

Usage: Performs specified SQL statement and informs caller of the results. The **actions** parameter returns a value of 0 if no additional action is required. If actions are required to complete the statement, one or more of the following bit flags are returned.

Flag	Action
-----	-----
SPQLDATA	Data is returned(see spqlfetch())
SPQLCOLINFO	Column information is returned(see spqlcolinfo())

Parameters:

void *contok

The connection used to execute the SQL statement.

char *stmtbuf

A buffer that holds the SQL statement to perform.

int stmtlen

The length of the SQL statement in buffer; -1 if null-terminated.

int *actions

Returns post-processing notification bit flags.

void **stmttok

Returns a statement token to use in post-processing the SQL statement results. See post-processing action definitions for use of statement token.

Returns: 0 if the SQL statement is successfully prepared/executed; SPQL_BADSTMT if the SQL statement specified in the statement buffer is prepared incorrectly; SPQL_NOMEM if **spqlperform** cannot allocate memory for the statement token.

[spqlfreestok\(\)](#)

Generates a free statement token from **spqlperform()**.

```
int spqlfreestok(void *stmttok);
```

Usage: Free resources used for the statement token from [spqlperform\(\)](#). Call **spqlfreestok()** after the data/information from the statement token has been extracted. You may call this function before all selected rows from the **spqlperform()** are read. If you do, the remaining unread rows (from the previous select) are discarded.

Parameters:

void *stmttok

Statement token to free

Returns: 0 if successful.

[spqltabinfo\(\)](#)

Gets table information from a statement token.

```
int spqltabinfo(void *stmttok, spqltinfo_t **tinfo)
```

Usage: Interrogates the statement token for table information. Upon return of the call, updates **tinfo**

with the pointer to the `spqltinfo_t` structure in the statement.

Note: Treat the structure accessed by the returned pointer as read-only memory.

Parameters:

void *stmttok

The statement token to use to access table information from a 'select'.

spqltinfo **tinfo

Returns pointer to `spqltinfo_t` structure into the statement token memory.

Returns: 0 for successful completion.

[spqlcolinfo\(\)](#)

Gets column information from a statement token.

```
int spqlcolinfo(void *stmttok, int *ncols, spqlcinfo_t **colvec)
```

Usage: Interrogates token for column information. Upon return of the call, updates `ncols` with the column count selected in the statement and updates `colvec` with the pointer to the vector of `spqlcol_t` structures in the statement.

Note: Treat structures accessed by the returned pointer as read-only memory.

Parameters:

void *stmttok

The statement token to use to access column information from 'select'.

int *ncols

Returns in the statement token the number of columns selected.

spqlcinfo **colvec

Returns in the statement token a pointer to the array of `spqlcinfo_t` structures.

Returns: 0 if successful.

[spqlfetch\(\)](#)

Gets row data from a statement token.

```
int spqlfetch(void *stmttok, void **bufptr, int *bufsize)
```

Usage: Fetches the rows returned from executing a statement. Each call to `spqlfetch` returns a row from a statement to the caller's buffer. If `bufptr` contains a NULL value, the routine returns a

pointer to a buffer containing the next row. If the value is not NULL, it assumes that the buffer is owned by the caller and returns the data to the caller's buffer. In either case, **bufsize** is updated with the row length returned. Callers that use locate-mode spqlfetch semantics (that is, who specify **bufptr** as NULL), should NEVER FREE the memory pointer returned by spqlfetch! A call to spqlfetch(), after all rows for the statement are returned, returns a **bufsize** of 0.

Parameters:

void *stmttok

The statement token to use to access row data from 'select'.

void **bufptr

Contains a pointer to the caller's row buffer to fill with row data. If it is NULL on entry, it returns a pointer to the internal statement buffer.

int *bufsize

Returns the size of the row buffer that was returned to the caller.

Returns: 0 if successful; SPQL_ENDDATA if the statement has no more rows to return; SPQL_FETCHFAILED if there is an unexpected failure while fetching the next row buffer.

[spqlgmsg\(\)](#)

Accesses thread-specific error/diagnostic message buffer contents.

```
int spqlgmsg(char **mbuf)
```

Usage: Returns a pointer to the threads error/diagnostic message buffer. Call **spqlgmsg()** to get any diagnostic messages if you encounter an error executing an SPQL function. If there is message information, spqlgmsg() returns the message pointer in the **mbuf** parameter as well as the length of the message (the function return value).

Parameters:

char **mbuf

Returns a pointer to the thread's error/diagnostic message buffer. If mbuf is NULL, there is no message information. The call also returns the length of the thread's error/diagnostic message buffer. A 0 indicates that no message exists.

[SPQL Function Return Codes](#)

Some SPQL functions generate return codes, allowing you to check the value and take appropriate action in your application code. Typically, the application action taken upon receiving an error code, is a call to spqlgmsg() to get the contents of the diagnostic buffer. The program can then display the buffer's contents to the user or write the contents to a log.

The following return codes are classified, in general, by their state: **positive** [(WARNING), (SUCCESS)] or **negative** [(ERROR)].

- [SPQL_SUCCESS\(==0\)](#)
- [SPQL_ENDDATA\(WARNING\)](#)
- [SPQL_INITFAILED\(ERROR\)](#)
- [SPQL_NOMEM](#)
- [SPQL_CONFAILED\(ERROR\)](#)
- [SPQL_BADSTMT\(ERROR\)](#)

[SPQL_SUCCESS\(==0\)](#)

Successful completion of the SPQL function call.

[SPQL_ENDDATA\(WARNING\)](#)

All rows selected were read from the statement token.

[SPQL_INITFAILED\(ERROR\)](#)

Initialization failure. (It is unsafe for your application to make additional SPQL calls if this error occurs.)

[SPQL_NOMEM](#)

Unable to allocate memory for some type of SPQL data structure. Check the diagnostic buffer for details.

[SPQL_CONFAILED\(ERROR\)](#)

Unable to make a connection to an SPD Server SQL server. Check the diagnostic buffer for details.

[SPQL_BADSTMT\(ERROR\)](#)

SQL statement is incorrectly formatted for submission to sqlprepare(). Either the statement is blank (all white space) or contains contiguous non-white space characters.

Optimizing SAS Scalable Performance Data (SPD) Server Performance

- [SAS Scalable Performance Data \(SPD\) Server Performance and Usage Tips](#)
- [Symmetric Multiple Processor \(SMP\) Utilization](#)
- [File System Performance Concepts](#)
- [LIBNAME Domains](#)
 - [Data and Index Separation](#)
 - [Configuring a LIBNAME Domain](#)
- [Loading Data into a SAS Scalable Performance Data \(SPD\) Server Host](#)
- [Loading Indexes in Parallel](#)
- [Truncating Tables](#)
- [Optimizing WHERE Clauses](#)
- [SAS Scalable Performance Data \(SPD\) Server Indexing](#)
- [WHERE Clause Planner](#)
- [How to Affect the WHERE Planner](#)
- [WHERE Clause Examples](#)
- [Server-Side Sorting](#)

[SAS Scalable Performance Data \(SPD\) Server Performance and Usage Tips](#)

SAS Scalable Performance Data (SPD) Server gives good performance when run using default configuration settings. To realize the full benefits of SAS Scalable Performance Data (SPD) Server's design and capabilities, you must configure some of the software's options to modify the default behaviors. The configuration changes will depend on the computing environment, table size and complexity, and indexing structures.

You use SAS/MACRO variables that are specific to SAS Scalable Performance Data (SPD) Server and SAS statement options (LIBNAME options and table options) to configure SAS Scalable Performance Data (SPD) Server for optimum performance.

[Symmetric Multiple Processor \(SMP\) Utilization](#)

A cornerstone of SAS Scalable Performance Data (SPD) Server's power is the ability to perform parallel processing. Parallel processing uses multiple processors to execute more than one set of instructions, or threads, concurrently. SAS Scalable Performance Data (SPD) Server is oriented to exploit parallelism whenever it can improve transaction times and processor utilization.

A fundamental question about parallelism is whether using additional CPUs on a specific problem will deliver data faster. Extra CPUs do not guarantee faster results every time. The amount of CPU-intensive work that a thread must do needs to last long enough to justify the cost of the thread. The cost of the thread is creating it, managing it, and interacting with other threads involved in the same parallel algorithm.

If not properly matched to the workload, the parallel algorithm can use more CPU time without reducing data delivery time. Additional threads can create conflicting demands for critical system resources such as physical memory. Excessive execution times can occur if too many threads attempt to access a large table at the same time, because many threads demand large amounts

of physical memory. Extreme resource constraints can result in slower overall processing.

SAS Scalable Performance Data (SPD) Server focuses on the following areas to speed overall processing using parallelism:

- User-definable parallel execution blocks for SQL pass-through statements
- Parallel aggregation for common summary functions when performing SELECT [...] GROUP BY statements
- WHERE Clause evaluation for indexed and non-indexed strategies
- Overlapped table and concurrent index updates when appending to tables
- Index creation when creating multiple indexes
- Optimize PROC SORT/BY-clauses
- Pipelined read-ahead when concurrently accessing multiple tables

[File System Performance Concepts](#)

SAS Scalable Performance Data (SPD) Server uses several file types in its data storage model. Data objects in SAS Scalable Performance Data (SPD) Server are composed of one or more component files. Each component file is itself a collection of one or more disk files. These are called the partitions of the component.

Component files create partitions when any of the following conditions is true:

- The current partition exceeds the user-specified [PARTSIZE=](#) value: Subsequent partitions are allocated in cyclical fashion across the set of directories that are specified in the [DATAPATH=](#) statement for the LIBNAME domain. Partitioning uses file-level striping to create PARTSIZE-sized files that complement disk-level striping that your operating system's volume manager software creates. SAS Scalable Performance Data (SPD) Server uses a default PARTSIZE= setting of 16 MB. PARTSIZE= determines a unit of work for parallel operations that require full table scans. Examples of parallel operations that require full table scans are WHERE Clause evaluation and SQL GROUP-BY summarization. Trade-offs are balancing increased numbers of files used to store the table versus the work savings realized through parallel partitions. Extra partitions means files are opened to process a table, but with fewer rows in each partition.
- The current partition exceeds the RLIMIT_FILESIZE value: In UNIX systems, RLIMIT_FILESIZE is a system parameter that defines the maximum size of a single disk file. In Windows, SAS Scalable Performance Data (SPD) Server uses a default RLIMIT_FILESIZE value of 2 GB.
- The current partition exceeds the space on the file system where it has been created.

-
- [Defining Directories](#)
 - [Disk Striping](#)
 - [RAID Levels](#)
 - [Transient Storage](#)

[Defining Directories](#)

SAS Scalable Performance Data (SPD) Server allows the user to define a set of directories that contain component files and their partitions. Normally, a single directory path is constrained by some volume limit for the file system, or the maximum amount of disk space that the operating system understands.

Most UNIX and Windows systems offer a volume manager utility. You can use volume manager utilities to create file systems (volumes) that are greater than the available space on a single disk. System administrators can use these utilities to create large, multi-gigabyte volumes. These volumes may be spread across a number of disk

partitions, or even span multiple disk devices. Volume manager utilities generally support creation of disk volumes that implement one of the common RAID (redundant arrays of inexpensive disks) configuration levels.

Disk Striping

A defining feature of all RAID levels is disk striping. Striping organizes the linear address space of a volume into pieces that are spread across a collection of disk drive partitions. For example, a user may configure a volume across two 1 GB partitions on separate disk drives A and B with a stripe size of 64K bytes. Stripe 0 lives on drive A, stripe 1 lives on drive B, stripe 2 lives on drive A, and so on.

By distributing the stripes of a volume across multiple disks it is possible to

- achieve parallelism at the disk I/O level
- use multiple kernel threads to drive a block of I/O.

This also reduces contention and data transfer latency for a large block I/O because the physical transfer can be split across multiple disk controllers and drives.

RAID Levels

The following is a brief summary of RAID levels relevant to SAS Scalable Performance Data (SPD) Server:

RAID-0

High performance with low availability. Physically losing a disk means data is lost. No redundancy exists to recover volume stripes on a failed disk.

RAID-1

Disk mirroring for high availability. Every block is duplicated on another mirror disk, sometimes referred to as shadowing. In the event one disk is lost, the mirror disk is still likely to be intact, preserving the data. RAID-1 can also improve read performance since a device driver has two potential sources for the same data. The system can choose the drive that has the least load/latency at a given point in time. The down side to RAID-1: it requires twice the number of disk drives as RAID-0 to store a given amount of data.

RAID-5

High performance and high availability at the expense of resources. An error correcting code (ECC) is generated for each stripe written to disk. The ECC distributes the data in each logical stripe across physical stripes in such a way that if a given disk in the volume is lost, data in the logical stripe can still be recovered from the remaining physical stripes. RAID-5's downside is resource utilization; RAID-5 requires extra CPU cycles and extra disk space to transform and manage data using the ECC model.

RAID-1+0

Many RAID systems offer a combination of RAID-1 (pure disk mirroring) and RAID-0 (striping) to provide both redundancy and I/O parallelism in a configuration known as RAID-1+0 (sometimes referred to as RAID-10). Advantages are the same as for RAID-1 and RAID-0. The only disadvantage is the requirement for twice as much disk as the pure RAID-0 solution. Generally, this configuration tends to be a top performer if you have the disk resources to pursue it.

Regardless of RAID level, disk volumes should be hardware striped when using the SAS Scalable Performance Data (SPD) Server software. This is a significant way to improve performance. Without hardware striping, I/O will bottleneck and constrain SAS Scalable Performance Data (SPD) Server performance.

Transient Storage

You should configure a RAID-0 volume for WORKPATH= storage for your SAS Scalable Performance Data (SPD) Server. When sizing this RAID-0 volume, keep in mind that the WORKPATH= that you set up for a given SAS Scalable Performance Data (SPD) Server host must be shared by all of its SQL and LIBNAME proxy processes that exist at a given point in time. The SAS Scalable Performance Data (SPD) Server Frequently Asked Questions ([FAQ](#)) is a good source of information on estimating [disk space requirements for WORKPATH=](#).

Consider using one or more RAID-0 volumes to locate the database domains that will support TEMP=YES LIBNAME assignments. This LIBNAME statement option creates a temporary storage domain that exists only for the duration of the LIBNAME assignment. This is the SAS Scalable Performance Data (SPD) Server equivalent of the SAS WORK library. All data objects (tables, catalogs, utility files) that are created in the TEMP=YES temporary domain are automatically deleted when you end the SAS session.

LIBNAME Domains

LIBNAME domains define the primary directory path and can optionally define other directories for placing the data and index components of SAS Scalable Performance Data (SPD) Server tables. The [METAPATH=](#), [DATAPATH=](#), and [INDEXPATH=](#) LIBNAME definition options determine the placement of SAS Scalable Performance Data (SPD) Server's component and partition files.

- [Data and Index Separation](#)
- [Configuring a LIBNAME Domain](#)

Data and Index Separation

The section on [File System Performance Concepts](#) discussed how distributing I/O load across different disk drives can improve performance. Further load distribution can be achieved by separating data and index components of SAS Scalable Performance Data (SPD) Server tables. To do this, use the [DATAPATH=](#) and [INDEXPATH=](#) options when configuring LIBNAME domains.

For example, when performing complex WHERE Clause evaluations, multiple threads are active on index component files and the data component file at the same time. Splitting the index and data file components onto different volumes can improve performance by reducing disk contention and increasing the level of parallelism down to the disk access level.

A word of caution when using DATAPATH= and INDEXPATH= options to distribute the data and index components: take extra care when performing and restoring disk backups of SAS Scalable Performance Data (SPD) Server tables using a system backup and restore utility. When making a backup, ensure that the [metadata, data, and index component partition files](#) are of the same generation and are in their respective directories.

When restoring a backup, restore the component partitions to the same directories where they were created. To avoid this restore problem, create symbolic links with the original directory path that point to the restore directories. Of course, if the components are not separated using the path options, this restore issue does not apply.

The backup and restore issues are not an issue when using the SAS Scalable Performance Data (SPD) Server Backup and Restore Utilities. These utilities resolve any component files when backing up or restoring tables. More information on SAS Scalable Performance Data (SPD) Server Backup and Restore Utilities is available in the

Configuring a LIBNAME Domain

Suppose a user has four volumes designated. Volumes exist for (1) SAS Scalable Performance Data (SPD) Server metadata, (2) data components, (3) index components, and (4) proxy working storage, as follows

- **/dmart_domain** is a 4 GB volume
- **/dmart_data** is a 40 GB volume
- **/dmart_index** is a 40 GB volume
- **/spds_work** is a 10 GB volume

The user wants to configure a LIBNAME domain called **dmart** to use **/dmart_domain** for the primary directory, with data components going to **/dmart_data**, and index components going to **/dmart_index**. The **/spds_work** volume should be configured for proxy working storage.

The configuration is made in two steps:

1. In the server parameter file (-parmfile) enter the following line:

```
WORKPATH=/spds_work;
```

2. In the SAS Scalable Performance Data (SPD) Server LIBNAME file (-libnamefile) enter the following domain definition:

```
libname=dmart  
  path=/dmart_domain  
  roptions="datapath=('/dmart_data')  
  indexpath=('/dmart_index')";
```

Loading Data into an SAS Scalable Performance Data (SPD) Server Host

SAS Scalable Performance Data (SPD) Server's emphasis on complete LIBNAME compatibility means that when you access SAS Scalable Performance Data (SPD) Server, the standard procedures used to create tables in SAS apply to SAS Scalable Performance Data (SPD) Server tables as well.

Using SAS, you can load data into SAS Scalable Performance Data (SPD) Server tables using DATA step programs, PROC COPY or PROC APPEND, and SCL applications. You also can use SQL pass-through to load SAS Scalable Performance Data (SPD) Server tables. The SAS Scalable Performance Data (SPD) Server SQL extensions for the LOAD TABLE and COPY TABLE statements provide further support.

Use LOAD TABLE to load a table from the projected columns of an SQL SELECT statement and create indexes, all in a single pass. LOAD TABLE exploits multi-thread table I/O and index creation. The multi-thread table I/O and index creation overlaps with the SELECT statement that extracts the data from its source tables.

Use COPY TABLE to copy an existing SAS Scalable Performance Data (SPD) Server table to a new table and include indexes as part of the copy operation. It offers the same parallel table and index I/O and overlapped input as the LOAD TABLE command.

The COPY TABLE and LOAD TABLE statements work only for source and target tables on the local machine.

Table Loading Techniques

The SAS data storage model adds rows to a data set one at a time. The SAS Scalable Performance Data (SPD) Server I/O engine buffers rows to be added from the SAS application and performs block adds using a highly efficient pipelined append protocol when communicating with the proxy.

- [Parallel Table Load Technique Using PROC APPEND](#)
- [Parallel Table Load Technique Using SQL Pass-Through](#)
- [Parallel Pass-Through Table Load and Data Subset](#)
- [Parallel Pass-Through Table Copy](#)

To achieve significant improvements in building a table, create the empty table first, defining indexes on the desired columns. Then, use PROC APPEND to populate the table and indexes. The example below demonstrates this technique.

Parallel Table Load Technique Using PROC APPEND

```
/* Create an empty SPD Server table with the same      */
/* columns and column attributes as the existing      */
/* SAS table.                                         */

data spdslib.cars;
set somelib.cars(obs=0);
run;

/* Create indexes for the empty table so the indexes */
/* are appended in parallel with the table appends. */

PROC DATASETS lib=spdslib;
  modify cars;
  index create make;
  index create origin;
  index create mpg;
quit;

/* PROC APPEND SAS table Cars to SPD Server table   */
/* Cars. The append to the SPD Server table and    */
/* its indexes will occur in parallel.              */

PROC APPEND
  base=spdslib.cars
  data=somelib.cars;
run;
```

If you are using SQL pass-through, consider using the LOAD TABLE command to perform the same operation. LOAD TABLE encapsulates the sequence of SAS DATA and PROC steps into an even more powerful technique for gaining maximum performance when loading a new table. The following example demonstrates the same table construction using LOAD TABLE and SQL pass-through:

Parallel Table Load Technique Using SQL Pass-Through

```
/* Create a copy of the SPD Server table Cars and */
/* its index from Example 1 to another SPD Server */
/* table carload using pass-through LOAD command. */
/* The table creation of the SPD Server table */
/* carload and its indexes will occur in parallel. */

execute(
load table carload with
  index make
    on (make),
  index origin
    on (origin),
  index mpg
    on (mpg)
  as select *
  from cars
) by sasspds;
```

Parallel Pass-Through Table Load and Data Subset

```
/* Create a subset of the SPD Server table Cars */
/* from Example 1 to another SPD Server table */
/* Fordcar using the pass-through LOAD command. */
/* The table creation of the SPD Server table */
/* Fordcar and its indexes occurs in parallel. */

execute(
load table fordcar with
  index origin
    on (origin),
  index mpg
    on (mpg)
  as select *
  from cars
  where make="ford"
) by sasspds;
```

Parallel Pass-Through Table Copy

```
/* Create a copy of the SPD Server table Cars and */
/* all its indexes from Example 1 to another Data */
/* Server table Copycars using the pass-through */
/* COPY command. The table creation of the Data */
/* Server table Copycars and its indexes will */
/* occur in parallel. */
```



```
execute(  
  copy table copycars  
    from cars  
  ) by sasspds;
```

Loading Indexes in Parallel

A significant strength of SAS Scalable Performance Data (SPD) Server is efficient creation, maintenance, and use of table indexes. Indexing can greatly speed the evaluation of WHERE Clause queries. The index can also be a source of sort order when performing BY-clause processing. The index is also used directly by some SAS applications. For example, PROC SQL uses indexes to efficiently evaluate equi-joins.

- [Parallel Index Creation](#)
- [Parallel Index Updates](#)

Parallel Index Creation

SAS Scalable Performance Data (SPD) Server supports parallel index creation using asynchronous index options. To enable asynchronous parallel index creation, either submit the [SPDSIASY=YES](#) macro variable prior to creating an index in SAS, or use the [ASYNCINDEX=YES](#) table option.

Both the macro variable and the table option apply to the DATA step INDEX= processing as well as to PROC DATASETS INDEX CREATE commands. Either method allows all of the declared indexes to be populated with a single scan of the table. A single scan is a substantial improvement over making multiple passes through the data to build each index serially.

As always, there is a price for parallelism. To create multiple indexes requires enough WORKPATH= disk space to create all of the key sorts at the same time. The PROC DATASETS structure has the flexibility to allow batched parallel index creation by using multiple MODIFY groups. The Parallel Index Creation example below inserts INDEX CREATE statements between two successive MODIFY statements resulting in a parallel creation group.

Parallel Index Creation Example

```
DATA foo.patient_info;  
  length  
    last_name $10  
    first_name $20  
    patient_class $2  
    patient_sex $1;  
  
  patient_no=10;  
    last_name="Doe";  
    first_name="John";  
    patient_class="XY";  
    patient_age=33;  
    patient_sex="M";  
  
run;
```

```

%let spdsiasy=YES;
PROC DATASETS lib=foo;
  modify patient_info;
    index create
      patient_no
      patient_class;
  modify patient_info;
    index create
      last_name
      first_name;
  modify patient_info;
    index create
      whole_name=(last_name first_name)
      class_sex=(patient_class patient_sex);
quit;

```

Indexes for PATIENT_NO and PATIENT_CLASS are created in parallel, indexes for LAST_NAME and FIRST_NAME are created in parallel, and indexes for WHOLE_NAME and CLASS_SEX are created in parallel.

Parallel Index Updates

SAS Scalable Performance Data (SPD) Server also supports parallel index updates during table append operations. Multiple threads enable overlap of data transfer to the proxy, as well as updates of the data store and index files. SAS Scalable Performance Data (SPD) Server decomposes table append operations into a set of steps that can be performed in parallel. The level of parallelism attained depends on the number of indexes that are present on the table. The more indexes you have, the greater the exploitation of parallelism during the append processing. As with parallel index creation, parallel index updates use WORKPATH= disk space for the key sorts that are part of the index append processing.

Truncating Tables

The Truncate command is a PROC SPDO command that allows the deletion of all rows in a table without deleting the table structure or metadata. The PROC SPDO truncate command is shaded for emphasis in the code example below.

```

%let host=kaboom ;
%let port=5191 ;
%let domain=path2 ;

libname &domain sasspds "&domain"
  server=&host.&port
  user='anonymous'
  ip=YES ;

/* create a table */

data &domain..staceys_table ;

  do i = 1 to 100 ;

```

```

    output ;
end ;
run ;

/* verify the contents of the created table */

PROC CONTENTS data=&domain..staceys_table ;
run ;

/* SPDO Truncate command deletes the table */
/* data but leaves the table structure in */
/* place so new data can be appended */

PROC SPDO lib=&domain ;
set acluser ;
Truncate staceys_table ;

quit ;

/* verify that no rows or data remain in */
/* the structure of staceys_table */

PROC CONTENTS data=&domain..staceys_table ;
run ;

```

Optimizing WHERE Clauses

SAS Scalable Performance Data (SPD) Server includes more advanced methods to optimize WHERE Clauses. Before SAS Scalable Performance Data (SPD) Server 4.0, the rule-based, heuristic WHERE Clause planner WHINIT was used to manually "tune" queries for performance. SAS Scalable Performance Data (SPD) Server provides dynamic WHERE Clause costing, an automatic feature which can replace the need to manually "tune" queries. SAS Scalable Performance Data (SPD) Server dynamic WHERE-costing uses factors of *duplicity* and *distribution* to calculate relative processor "costs" of various WHERE Clause options. SAS Scalable Performance Data (SPD) Server users can use server parameter commands in the `spdsserv.parm` file or macro variables to turn dynamic WHERE-costing on and off. If dynamic WHERE-costing is turned off, SAS Scalable Performance Data (SPD) Server reverts to using the rules-based WHERE Clause planner.

WHERE Clause Definitions and Terminology

- **WHERE Clauses** are selection criteria for a query that specify one or more boolean predicates. Implementing the criteria, SAS Scalable Performance Data (SPD) Server selects only records that *satisfy* the WHERE clause.
- **Predicates** are the building blocks of WHERE clauses. Use them stand-alone or combine them with the operators AND and/or OR to form complex WHERE clauses. An example of a WHERE Clause is

```
"where x > 1 and y in (1 2 3)"
```

In this example, there are two predicates, "`x > 1`" and "`y in (1 2 3)`". You specify the negative of a predicate by using "not". For example, "`where x > 1 and not (y in (1 2 3))`".

- **Boolean logic** determines whether two predicates, joined with an AND or OR, are true (satisfies) or false (does not satisfy) the specification. The AND operator requires that all predicates be true for the entire expression to be true. For example, the expression "p1 AND p2 AND p3", is true only if all three predicates (p1, p2 and p3) are true. In contrast, the OR operator requires only one predicate to be true for the entire expression to be true.

For the WHERE clause "(x < 5 or y in (1 2 3)) and z = 10", the following truth table describes the overall result (truth):

"x < 5 ?"	"y in (1 2 3) ?"	"z = 10 ?"	Result
=====	=====	=====	=====
False	False	False	False
False	False	True	False
False	True	False	False
False	True	True	True
True	False	False	False
True	False	True	True
True	True	False	False
True	True	True	True

- **Indexes** are structures associated with tables that permit SAS Scalable Performance Data (SPD) Server to quickly access records that satisfy an *indexed* predicate. In an example WHERE clause, "where x = 10 and y > 11", SAS Scalable Performance Data (SPD) Server selects the best index on column "x" to directly retrieve records that have a value of 10 in the "x" column. If no index exists for "x", SAS Scalable Performance Data (SPD) Server must sequentially read each record in the table searching for "x" equal to 10.
- **Simple and composite indexes:** Simple indexes index a single column; composite indexes index two or more columns. The list of column(s) in an index is sometimes called the index key.
- **Parallelism** is the SAS Scalable Performance Data (SPD) Server capability that enables multiple threads to execute in parallel. Using multiple processors in parallel mode is sometimes called 'divide and conquer' processing. SAS Scalable Performance Data (SPD) Server uses parallelism to evaluate the multiple indexes that are involved in more complicated WHERE clauses.

SAS Scalable Performance Data (SPD) Server Indexing

SAS Scalable Performance Data (SPD) Server tables may have one or more indexes. There are a combination of four different indexing strategies a table can use, and the choice depends on the data populating the table, the size of the table, and the types of queries that will be executed against the table.

SAS Scalable Performance Data (SPD) Server indexing evaluates the processor "cost" of a WHERE Clause. The section [Costing Using Duplicity and Distribution Values](#) shows how factors of *duplicity* and *distribution* are used to choose the evaluation strategy that will perform the WHERE Clause at the smallest processor "cost". The five evaluation strategies that the WHERE Clause planner uses are EVAL 1, EVAL 2, EVAL 3, EVAL 4, and EVAL 5. The differing EVAL strategies calculate the number of rows that will be required to execute a given query.

"True" rows are rows that contain the variable values specified in a WHERE Clause. "False" rows do not contain the variable value specified in the clause. EVAL 1, EVAL 3, EVAL 4, and EVAL 5 evaluate "true" rows in the table using indices. EVAL 2 evaluates true rows of a table without using indices. EVAL strategies are explored in more detail in the section below on [WHERE Clause EVAL Strategies](#).

- [SPD Indexes](#)
- [MINMAX Indexes](#)

SPD Indexes

SAS Scalable Performance Data (SPD) Server uses segmented indices. A segmented index is created by dividing the index of a table into equally sized ranges of rows. Each range of rows is called a segment, or slot. You use the `SEGSIZE=` setting to define the size of the segment. A series of sub-indices each point to blocks of rows in the table. By default, SAS Scalable Performance Data (SPD) Server creates an index segment for every 8192 rows in a table.

The SPD segmented index facilitates SAS Scalable Performance Data (SPD) Server's parallel evaluation of WHERE Clauses with an indexed predicate. First, the SPD index supports a pre-evaluation phase to determine which segments contain values that satisfy the predicate. Pre-evaluation speeds queries by eliminating segments that do not contain any possible values. Then, up to [SPDSTCNT=](#) threads are launched to query the remaining index segments. The threads query the segments of the SPD index in parallel to retrieve the segment rows that satisfy the predicate. When all segments have been queried, the per-segment results are accumulated to determine the rows that satisfy the predicate. If the query contains multiple indexed predicates, then those predicates are also evaluated in parallel. When all predicates have been completed, their results are accumulated to determine the rows that satisfy the query.

MINMAX Indexes

SAS Scalable Performance Data (SPD) Server contains a new table option called `MINMAXVARLIST=`. The primary purpose of the `MINMAXVARLIST=` table option is for use with SAS Scalable Performance Data (SPD) Server [dynamic cluster tables](#) where specific members in the dynamic cluster contain a set or range of values, such as sales data for a given month. When a SAS Scalable Performance Data (SPD) Server SQL subsetting WHERE Clause specifies specific months from a range of sales data, the WHERE planner checks the min/max indexes. Based on the min/max index information, the SAS Scalable Performance Data (SPD) Server WHERE planner includes or eliminates member tables in the dynamic cluster for evaluation.

Use the `MINMAXVARLIST=` table option with either numeric or character-based columns. `MINMAXVARLIST=` uses the list of columns you submit to build an index. The `MINMAXVARLIST=` index contains only the minimum and maximum values for each column. The WHERE Clause planner uses the index to filter SQL predicates quickly, and to include or eliminate member tables belonging to the cluster table from the evaluation.

Although the `MINMAXVARLIST=` table option is primarily intended for use with dynamic clusters, it also works on standard SAS Scalable Performance Data (SPD) Server tables. `MINMAXVARLIST=` can help reduce the need to create many indexes on a table, which can save valuable resources and space.

The `MINMAXVARLIST=` table option is only available when a table is being created or defined. If a table has a `MINMAXVARLIST=` type of index, moving or copying the table will destroy the index unless `MINMAXVARLIST=` is specified in the table output.

```
%let domain=path3 ;
%let host=kaboom ;
%let port=5201 ;

libname &domain sasspds "&domain"
    server=&host..&port
    user='anonymous' ;
```

```

/* Create three tables called */
/* xy1, xy2, and xy3.      */

data &domain..xy1(minmaxvarlist=(x y));
  do x = 1 to 10;
  do y = 1 to 3;
  output;
  end;
end;
run;

data &domain..xy2(minmaxvarlist=(x y));
  do x = 11 to 20;
  do y = 4 to 6 ;
  output;
  end;
end;
run;

data &domain..xy3(minmaxvarlist=(x y));
  do x = 21 to 30;
  do y = 7 to 9 ;
  output;
  end;
end;
run;

/* Create a dynamic cluster table */
/* called cluster_table out of    */
/* new tables xy1, xy2, and xy3   */

PROC SPDO library=&domain ;
  cluster create cluster_table
    mem=xy1
    mem=xy2
    mem=xy3
    maxslot=10;
quit;

/* Enable WHERE evaluation to see */
/* how the SQL planner selects    */
/* members from the cluster. Each */
/* member is evaluated using the  */
/* min-max index.                */

%let SPDSWDEB=YES;

/* The first member has true rows */

PROC PRINT data=&domain..cluster_table ;
  where x eq 3 and y eq 3;
run;

```

```

/* Examine the other tables */

PROC PRINT data=&domain..cluster_table ;
  where x eq 19
  and y eq 4 ;
run;

PROC PRINT data=&domain..cluster_table ;
  where x eq 22
  and y eq 9;
run;

PROC PRINT data=&domain..cluster_table ;
  where x between 1 and 10
  and y eq 3;
run;

PROC PRINT data=&domain..cluster_table ;
  where x between 11 and 30
  and y eq 8 ;
run;

/* Delete the dynamic cluster table. */

PROC SPDO library=&domain ;
  cluster undo cluster_table ;
quit;

PROC DATASETS lib=&domain nolist;
  delete xy1 xy2 xy3 ;
quit ;

```

WHERE Clause Planner

The WHERE Clause Planner implemented in SAS Scalable Performance Data (SPD) Server avoids computation-intensive operations and uses simple computations where possible. WHERE Clauses in large database operations can be very resource-intensive operations. In SAS Scalable Performance Data (SPD) Server 3.x and earlier releases, query authors often needed to manually "tune" queries for performance. The "tuning" was accomplished using macro variables and index settings. The WHERE Clause planner integrated into SAS Scalable Performance Data (SPD) Server does the "tuning" work for the user by automatically costing the different approaches to index evaluation.

- [WHERE-Costing Using Duplicity and Distribution Values](#)
- [WHINIT: Indexed and Non-Indexed Predicates](#)
- [WHERE Clause EVAL Strategies](#)
- [Assigning EVAL Strategies](#)
- [Sample WHINIT Output](#)
- [WHINIT Output Return Keywords](#)

- o [Composite Index Permutations](#)

WHERE-Costing Using Duplicity and Distribution Values

Two key factors are used to evaluate, or "cost" WHERE Clause indices. The factors are *duplicity* and *distribution*.

Duplicity refers to the proportion expressed by the number of rows in a table divided by the number of distinct values in the index. When many observations in a table hold the same value for a given variable, the variable value is said to have a high duplicity. An example of a table with high duplicity might be a table of unleaded gasoline prices from service stations in the same area of a large city.

Conversely, when a table has only one or few observations that contain a given value for a variable, then that value can be described as low duplicity. An example of a table with low duplicity might be an office phone directory, where the variable for phone extension is always unique.

The duplicity value for an index ranges from 1 to the number of rows in the table. Indices with a duplicity value of 1 are unique. Indices with high duplicity generate a score that is close to the number of rows in the table.

Distribution refers to the sequential proximity between observations for values of a variable that are repeated throughout the variable's data set distribution. When a certain value for a variable exists in many observations that are scattered uniformly throughout the table, that value is said to have a wide distribution. If a variable value exists in many contiguous or nearly contiguous rows, the distribution is clustered.

WHERE Clause EVAL Strategies

SAS Scalable Performance Data (SPD) Server indexing keeps track of the duplicity and distribution of variable values in a table and uses them to calculate the cost of a WHERE Clause. The WHERE Clause planner uses four evaluation strategies to determine the number of rows that will be required to execute a given query. The four evaluation strategies are EVAL 1, EVAL 2, EVAL 3, and EVAL 4. "True" rows are rows that contain the variable values specified in a WHERE Clause. "False" rows do not contain the variable value specified in the clause.

EVAL 1, EVAL 3, EVAL 4, and EVAL 5 evaluate "true" rows in the table using indices. EVAL 2 evaluates true rows of a table without using indices.

- **EVAL 1** evaluates "true" rows using an index to locate the true rows in each segment of the table. The index evaluation process generates a list of row IDs per segment. EVAL 1 accepts WHERE Clause operators for equivalency expressions such as **EQ**, **=**, **LE**, **<=**, **LT**, **<**, **GE**, **>=**, **GT**, **>**, **IN**, and **BETWEEN**. EVAL 1 uses threaded parallel processing across the index segments to permit concurrent evaluation of multiple indices. EVAL 1 combines multiple segment bitmaps from queries that use multiple indices to generate the list of row IDs per segment.
- **EVAL 2** takes "true" rows as determined by EVAL 1, EVAL 3, or EVAL 4, then uses brute force to eliminate any rows shown to be "false", leaving a table which contains only "true" rows. EVAL 2 processes all rows of a table when no index evaluation is possible. For example, no index evaluation is possible when an index is not present or when some predecessor function performs an operation that invalidates the index.
- **EVAL 3** is a single index sequential process. Use EVAL 3 when the number of rows returned by an index is unique or nearly unique (when duplicity is low). EVAL 3 returns a list of "true" rows for the entire table. EVAL 3 only supports the equality operators **EQ** and **=**.

- **EVAL 4** is similar to EVAL 3 but supports a larger set of inequality and inclusion operators, such as **IN**, **GT**, **GE**, **LT**, **LE**, and **BETWEEN**.
- **EVAL 5** can operate when the SAS Scalable Performance Data (SPD) Server Index Scan Facility is used. The EVAL 5 strategy uses index metadata and aggregate SQL functions to evaluate "true" rows. The EVAL 5 strategy does not require a table scan.

For example, when *x* is indexed, and SAS Scalable Performance Data (SPD) Server uses EVAL 5 to evaluate the SQL expression

```
count(*) where x=5 ,
```

the index metadata is scanned for the condition, "*x* = 5" instead of performing table scans. The EVAL 5 strategy supports the min(), max(), count(), count(distinct), nmiss(), and range() functions. The EVAL 5 strategy cannot be used on SQL expressions which uses functions other than those listed above.

The WHERE Clause planner in SAS Scalable Performance Data (SPD) Server 3.x relied heavily on EVAL 1 and EVAL 2 threaded strategies to evaluate most clauses. Sometimes the SAS Scalable Performance Data (SPD) Server 3.x EVAL 1 and EVAL 2 strategies would over-thread and over-manipulate indices during the evaluations during WHERE Clause evaluation. This resulted in reduced performance or excessive resource consumption. With SAS Scalable Performance Data (SPD) Server 4.4's WHERE Clause costing in place, EVAL 3 and EVAL 4 strategies are more suitable evaluation engines which conserve resources and boost processor performance.

Assigning EVAL Strategies

The SAS Scalable Performance Data (SPD) Server WHERE Clause planner uses the following logic when selecting an EVAL strategy to evaluate expressions:

When the planner encounters a WHERE Clause, it builds a tree that represents all of the possible predicate expressions. The objective of the WHERE Clause planner is to divide the set of predicate expressions into two trees. One tree collects predicate expressions which lack usable indices and are constrained to EVAL 2 evaluation. The remaining predicate expressions are put in the other tree. Each of the predicate expressions in the second tree are scanned and assigned an evaluation strategy of EVAL 1, EVAL 3, or EVAL 4, depending on the WHERE Clause costing values and the syntax used in the predicate expression .

The second tree, which does not use the EVAL 2 method, is scanned for predicate expressions that return values with high duplicity . When high duplicity predicate expressions are identified, they are ranked. The predicate expression with the highest duplicity value is set aside for an index-based evaluation. All of the other remaining predicate expressions are evaluated using the EVAL 2 tree strategy. The lowest duplicity predicate expression is evaluated using either the EVAL 3 or the EVAL 4 strategy. The syntax used in the predicate expression determines which of the two strategies to use. Frequently, the single index EVAL 3 or EVAL 4 is chosen because single index evaluations require smaller processing loads and yield reliable results. With a low processor overhead and a high data yield, there is no reason to include other indices when a single index is sufficient.

When the WHERE Clause planner determines that no predicate expressions meet the high duplicity criteria, it chooses the EVAL 1 strategy. Before the EVAL 1 operation is performed, the costing algorithm is run on the remaining predicates in order to prune any predicate expressions which represent large processor loads and large data yields. Predicate expressions which will require large processor loads and produce large data yields are moved to the EVAL 2 tree.

- o [Index Scan Facility](#)
- o [High Yield Predicate Expressions](#)
- o [High Processing Load Predicate Expressions](#)
- o [High Yield and High Processing Load Predicate Expressions](#)
- o [Turning WHERE Clause Costing Off](#)

Index Scan Facility

When SAS Scalable Performance Data (SPD) Server invokes the Index Scan Facility, and the SQL aggregate uses the specified supported functions for EVAL 5, the EVAL 5 strategy uses a fast index metadata scan to select SQL statements that meet the aggregate function criterion.

High Yield Predicate Expressions

A large, or high data yield expression has a high percentage of rows containing true segments. The default threshold for a high yield expression is one where less than 25% of the rows evaluated are returned by the predicate. At this point, processor costs related to index use begin increasing without proportional returns on the evaluation results.

High Processing Load Predicate Expressions

Predicate expressions that require high processing loads are predicates that usually require large amounts of index manipulation before they can complete. When the amount of index work that is required exceeds the work that is required to use an EVAL 2 strategy, the predicate expression will be best evaluated by the EVAL 2 tree. Open-ended predicate expressions that contain many syntax inequality operators such as GT and LT or many variations in syntax are good high work candidates for EVAL 2. High work predicate expressions are detected by comparing the number of unique values in the predicate expression to the number of unique values contained in the index.

High Yield and High Processing Load Predicate Expressions

When all predicate expressions in EVAL 1 are high yield or high processor load, SAS Scalable Performance Data (SPD) Server uses segmented costing. In segmented costing, "true" segments are passed to EVAL 2 for processing. EVAL 2 only processes table segments that can provide "true" rows for the WHERE Clause.

Turning WHERE Clause Costing Off

You can use the SAS Scalable Performance Data (SPD) Server `spdsserv.parm` parameter file to configure the default WHERECOSTING parameter setting to ON. If you want to turn WHERE Clause costing off within the scope of a job, you can use macros or a DATA step to turn WHERE Clause costing off and on:

- The SPDSWCST=NO macro setting turns off WHERE Clause costing.
- The SPDSWSEQ=YES macro overrides WHERE Clause costing and allows you to force a global EVAL3 or EVAL4 strategy.
- The WHERECOSTING parameter can be removed or set to NOWHERECOSTING in the `spdsserv.parm` file if you want to turn off costing for the entire server.

If you turn WHERE Clause costing off in the `spdsserv.parm` parameter file, or if you use the macro setting SPDSWCST=NO, the WHERE Clause planner reverts to the rules-based WHERE Clause planning of earlier versions of SAS Scalable Performance Data (SPD) Server.

WHINIT: Indexed and Non-Indexed Predicates

If SAS Scalable Performance Data (SPD) Server is not configured to use dynamic WHERE-costing, the WHERE Clause planner reverts to the rule-based heuristics of WHINIT. WHINIT uses rules to select indexes for the predicates, and then select the most appropriate EVAL strategy for the query.

WHINIT splits the WHERE clause, represented as a tree, into non-indexed and indexed parts. Non-indexed predicates include

- non-indexed columns
- functions
- columns that have indexes that WHINIT cannot use.

If the WHERE Clause planner places indexed predicates in the non-indexed tree, it is usually because the predicates involve an OR expression. An example of a predicate with an OR expression is, "where x = 1 or y = 2". Even if column "x" is indexed, WHINIT cannot use the index because the OR is disjunctive. As a result of the disjunctive OR, the planner cannot use the index, and places both the predicates, "x = 1" and "y = 2", into the non-indexed part of the WHERE tree.

Sample WHINIT Output

SAS users can use an SAS Scalable Performance Data (SPD) Server macro variable to view WHERE Clause planner output:

```
%let SPDSWDEB=YES;
```

The following is what the WHINIT plan might give for the following scenario:

- a WHERE clause of "**where a = 1 and b in (1 2 3) and d = 3 and (d + 3 = c)**"
- an SPD index IDX_ABC on columns (A B C)
- an SPD index D on column (D).

Note: The line numbers are for reference; they are NOT part of the actual output.

```
1:whinit: WHERE ((A=1) and B in (1, 2, 3) and (D=3) and (C=(D+3)))  
2:whinit: wh-tree presented
```

```

3:      /-NAME = [A]
4:      /-CEQ----|
5:      |
6:      \-LITN = [1]
7:      --LAND---|
8:      /-NAME = [B]
9:      |--IN-----|
10:     | /-LITN = [1]
11:     \-SET----|
12:     |--LITN = [2]
13:     \-LITN = [3]
14:     /-NAME = [D]
15:     |--CEQ----|
16:     \-LITN = [3]
17:     /-NAME = [C]
18:     \-CEQ----|
19:     | /-NAME = [D]
20:     \-AADD---|
21:     \-LITN = [3]
21:whinit: wh-tree after split
22:     /-NAME = [C]
23: --CEQ----|
24:     /-NAME = [D]
25:     \-AADD---|
26:     \-LITN = [3]
27:whinit: SBM-INDEX D uses 50% of segs (WITHIN maxsegratio 75%)
28:whinit: INDEX tree after split
29:     /-NAME = [A] <1>SBM-INDEX IDX_ABC (A,B)
30:     /-CEQ----|
31:     |
32:     \-LITN = [1]
33: --LAND---|
34:     /-NAME = [B]
35:     |--IN-----|
36:     | /-LITN = [1]
37:     \-SET----|

```

```

38:          |--LITN = [2]
          |
39:          \-LITN = [3]
          |
40:  /-NAME = [D] <2>SBM-INDEX D (D)
41:          \-CEQ----|
42:          \-LITN = [3]
42:whinit returns: ALL EVAL1(w/SEGLIST) EVAL2

```

Line 1 shows what the WHINIT Planner received. Do not be surprised -- what the Planner receives can differ from your entries. Sometimes SAS optimizes or transforms a WHERE clause before passing it to SAS Scalable Performance Data (SPD) Server. For example, it can eliminate entities such as NOTs, the union of set lists, and so on.

Lines 2 to 20 show the presented WHERE clause in a tree format. The tree format is a user-readable form of the actual WHERE clause that is processed by the SAS Scalable Performance Data (SPD) Server engine.

Lines 21 to 26 show the non-indexed WHERE tree, the result of splitting off the *indexed* part. The non-indexed WHERE tree can be "empty" or it can look the same as lines 2 to 20 if no indexes are selected. Bear in mind that it is the non-indexed part of the WHERE clause that WHINIT uses to filter records obtained by the indexed strategies (EVAL1, 3 or 4).

Lines 27 to 41 shows that the percentage of segments containing values selected from column D is with the maximum allowed to proceed with pre-segment logic. Therefore, only those segments that contain values that satisfy the where clause for column D will be included in further query processing for that column. Composite index IDX_ABC and simple index D are used to resolve the indexed WHERE clause predicates.

Line 42, the last line in our output, shows which strategies are used. The first keyword "ALL" indicates that SAS Scalable Performance Data (SPD) Server can identify correctly ALL resulting records, without help from the SAS System. First, SAS Scalable Performance Data (SPD) Server will call EVAL1, an indexed method, to quickly access a list of records which satisfy "**where a = 1 and b in (1 2 3)**" and "**d = 3**", then it will use EVAL2 to determine if "**c = d + 3**" is true on these records.

When output from EVAL1 displays the suffix "w/ seglist", as it does in the above output, it means that SPD indexes were detected, and that the indexes were used to filter out only the segments that satisfy the given indexed predicates. When EVAL1 has no suffix, it means that ALL segments will be evaluated.

SAS Scalable Performance Data (SPD) Server stores the minimum and maximum values for a table index in a global structure. WHINIT can use the numeric range to 'prune' predicates when the table index values are out of the min / max range. WHINIT output keywords can indicate pruning activity. For example, if WHINIT had determined that the values for "D" (in our WHERE clause) are between 5 and 13, then as a consequence, the predicate "**where d = 3**" could never be true. In this case, WHINIT would have pruned this predicate since it is logically impossible, or FALSE. Pruning can also affect higher nodes. If the "**d = 3**" predicate were deemed FALSE, then the AND sub tree would also be FALSE and would also have been pruned.

WHINIT Output Return Keywords

In the last line of the output, "ALL" is one of the following keywords that the Planner can display:

- **ALL** - SAS Scalable Performance Data (SPD) Server can evaluate ALL of the WHERE clause when

determining which records satisfy the clause.

- **SOME** - SAS Scalable Performance Data (SPD) Server can handle SOME, or part, of the WHERE clause; it will then need some of the SAS System to help identify resulting records.
- **NONE** - SAS Scalable Performance Data (SPD) Server cannot evaluate this WHERE clause; the SAS System will perform all evaluations.
- **TRUE** - SAS Scalable Performance Data (SPD) Server has determined that the entire WHERE clause is TRUE, and that all the records satisfy the given WHERE clause. (It did index minimum / maximum values or other checks , for example, "where 1 = 1".)
- **FALSE** - SAS Scalable Performance Data (SPD) Server determined that the WHERE clause is FALSE, that is, no records can satisfy the WHERE clause.
- **RC=number** - An internal error has occurred; the error number is displayed.
- **EVALx** - the EVAL strategies the Planner will use, "x" can be 1, 2, 3 or 4.

Composite Index Permutations

A composite index can involve one or more "in set" equality predicates, such as an index on columns (a b c). When WHINIT is presented with a WHERE Clause that has such a composite index, for example, "where a = 1 and b in (1 2 3) and c in (4 5)", it will generate all permutations of this compound key, probing the index for each value. In our example, six values are generated:

(a b c) = (1 1 4) (1 1 5) (1 2 4) (1 2 5) (1 3 4) (1 3 5)

The permutations start at the "back" end of the key to take advantage of locality: to locate keys with close values which access the same disk page. This means less input/output operations on the index.

How to Affect the WHERE Planner

- [Macro Variable: SPDSWCOST=](#)
- [Macro Variable: SPDSWDEB=](#)
- [Macro Variable: SPDSIRAT=](#)
- [Macro Variable SPDSNIDX = or Table Option NOINDEX=](#)
- [Macro Variable SPDSWSEQ=](#)
- [Server Parameter Option WHERECOSTING](#)
- [WERENOINDEX= Option](#)
- [When and Why Should I Suppress Indexes?](#)
- [Identical Parallel WHERE Clause Subsetting Results](#)

Macro Variable SPDSWCST=

To turn off dynamic WHERE-costing, specify

```
%let SPDSWCST=NO;
```

Macro Variable SPDSWDEB=

To turn on WHINIT planning output, specify

```
%let SPDSWDEB=YES;
```

Macro Variable SPDSIRAT=

To affect the WHERE-planner SPD index pre-evaluation, specify

```
%let SPDSIRAT=index-segment-ratio;
```

The SPDSIRAT= macro variable specifies a maximum percentage (ratio) for the number of segments in the hybrid bitmap which must contain the index value before the WHERE-planner should pre-evaluate a segment list.

The segment list enables the planner to launch threads only for segments that contain the value. If the value number exceeds the ratio, the planner performs no pre-evaluation. Instead, the planner launches a thread for each segment in the table.

The SPDSIRAT= macro variable option can be used to ensure that time spent in pre-evaluation does not exceed the cost of launching a thread for each segment in the table. By default SPDSIRAT= is set to 75 percent. This means that if an index value is contained in 75 percent or less of the index segments, the hybrid bitmap logic will pre-evaluate the value and return a list of segments to the WHERE Clause planner. If more than 75 percent of the index segments contain the target index value, the time spent on pre-evaluation might be more than the time saved by skipping a small number of segments.

For some tables 75 percent **may not** be the optimal setting. To determine a better setting, run a performance benchmark, adjust the percentage, and rerun the performance benchmark. Comparing results will show you how the specific data population you are querying responds to shifting the index-segment ratio. The allowable range to adjust the setting value is from 0 to 100, where 0 means **never** perform WHERE Clause pre-evaluation, and 100 means **always** perform WHERE Clause pre-evaluation.

Macro Variable SPDSNIDX= or Table Option NOINDEX=

To suppress WHINIT use of any index, specify the no index SAS Scalable Performance Data (SPD) Server macro variable or the corresponding SAS Scalable Performance Data (SPD) Server table option:

```
%let SPDSNIDX=YES;  
  
data _null_;  
set foo.a (noindex=yes);
```

Macro Variable SPDSWSEQ=

By default, when WHINIT detects equality predicates that have indexes, it chooses EVAL1. However, the user may decide that sequential EVAL3 or EVAL4 methods are better. For example, in an equality WHERE predicate such as "where x = 3", WHINIT will default EVAL1 to evaluate the clause. If a user knows that the table queried has only a few records that can satisfy this predicate, EVAL3 may be a better choice. To force WHINIT to choose EVAL3/4, specify:

```
%let SPDSWSEQ=YES;
```

Server Parameter Option [NO]WHERECOSTING

Controls whether the server uses dynamic WHERE-costing. When dynamic WHERE-costing is disabled, the rules-based WHINIT heuristic is used to "tune" WHERE Clauses for performance. The default setting is for NOWHERECOSTING.

WHEREINDEX Option

A user may decide that one or more indexes selected by a WHINIT plan are not the best choice. This can occur because WHINIT is rule-based, not cost-based. Sometimes WHINIT selects a less-than-optimal plan. WHINIT's use of specific indexes can be affected by specifying the SAS Scalable Performance Data (SPD) Server option [WHEREINDEX=](#) in your data step.

```
data _null_;  
  set foo.a (wherenoindex=(idx_abc d))
```

This example specifies that WHINIT not use index "idx_abc" and index "d".

When and Why Should I Suppress Indexes?

Most rule-based planners, including WHINIT from SAS Scalable Performance Data (SPD) Server, assume that the index has a uniform distribution of values between the upper and lower value boundaries. This means if data values range between 2 and 10, that there are an equal number of 3's and 4's, and so on. When the assumption of a uniform distribution is false, an indexed predicate can return a large number of records. In turn, this causes WHINIT's indexed plan to run slower than a sequential read of the entire table. In this case the index should be suppressed.

Here is another, more subtle instance. When the WHERE clause uses only the front part of the key, WHINIT selects a composite index. Assume an index **abcd** on columns A, B, C and D and an index **e** on column E, and specify the WHERE clause

```
where a = 3 and e = 5;
```

Normally, WHINIT will select both indexes (**abcd** and **e**) and choose EVAL1. However, using the index **abcd** just to interrogate **a** might return a large number of records. In this case, suppressing the **abcd** index may be a good idea. If so, WHINIT will still choose EVAL1 for **e = 5**, or EVAL3 if SPDSWEV1=NO, and EVAL2, the post-

filter, for **a = 3**.

Identical Parallel WHERE Clause Subsetting Results

Under certain circumstances, it is possible to perform parallel WHERE Clause subsetting on a table more than once and to receive slightly different results. This event can occur when submitting parallel WHERE Clause code to SAS Scalable Performance Data (SPD) Server that utilizes the SAS **OBS=nnnn** data set option.

The SAS **OBS=nnnn** data set option causes processing to end with the specified (nth) observation in a table. Because parallel WHERE Clause processing is threaded, subsetting a table and using the **OBS=nnnn** may not produce identical results from run to run, or different batch jobs using the same WHERE Clause code may produce slightly different results.

When a parallel WHERE-clause evaluation is split into multiple threads, SAS Scalable Performance Data (SPD) Server employs a multi-threading model that is designed to return rows as fast as possible. Some threads may be able to complete row scans incrementally faster than other threads, due to uneven loads across multiple processors or system contention issues. This inequity can create minute variances which can generate non-identical results to the same subsetting request.

If you have code that performs parallel WHERE Clause subsetting in conjunction with the **OBS=nnnn** data processing option, and if it is critical that successive WHERE Clause subsets on the same data must be identical, you can eliminate thread contention error by setting the thread count value for that operation to 1.

To set the SAS Scalable Performance Data (SPD) Server thread count value, you can use the [SPDSTCNT=](#) macro:

```
%let SPDSTCNT=1;
```

The same potential for subsetting variation applies when a data step uses the **OBS=nnnn** data processing option with a parallel by-clause, such as:

```
data test1;  
  set spds44.testdata (obs=1000);  
  where j in (1,5,25);  
  by i;  
run;
```

Use the SPDSTCNT= macro solution to ensure identical results across multiple identical table subsetting requests.

- [WHERE Clause Subsetting Variation Example](#)

WHERE Clause Subsetting Variation Example:

Job 1 and Job 2 use the same tables and data requests but produce non-identical results as seen in the respective Job 1 and Job 2 outputs.

To eliminate variation in the output, simply add the thread count statement

```
%let SPDSTCNT=1;
```

to the beginning of each job.

Job 1

```
data test1;
  set spds44.testdata
    (obs=1000);
  where j in (1,5,25);
run;

PROC SORT data=test1;
  by i;
run;

PROC PRINT data=test1
  (obs=10);
run;
```

Job 1 Output:

The SAS System 11:44 Monday, May 9, 2005 1

Obs	a	i	j	k
1		24601	1	1
2		24605	5	5
3		24625	25	0
4		24701	1	1
5		24705	5	5
6		24725	25	0
7		24801	1	1
8		24805	5	5
9		24825	25	0
10		24901	1	1

Job 2:

```
data test2;
  set spds44.testdata
    (obs=1000);
  where j in (1,5,25);
run;

PROC SORT data=test2;
  by i;
run;
```

```
PROC PRINT data=test2
  (obs=10);
run;
```

Job 2 Output:

The SAS System
11:44 Monday, May 9, 2005 1

Obs	a	i	j	k
1		1	1	1
2		5	5	5
3		25	25	0
4		101	1	1
5		105	5	5
6		125	25	0
7		201	1	1
8		205	5	5
9		225	25	0
10		301	1	1

WHERE Clause Examples

- [Example 1](#): "where i = 1 and j = 2 and m = 4"
- [Example 2](#): "where i in (1, 2, 3) and j in (4, 5, 6, 7) and k > 8 and m = 2"
- [Example 3](#): "where i = 1 and j > 5 and mod(k, 3) = 2"
- [Example 4](#): "where i = 1 and j > 5 and mod(k, 3) = 2" (the index IJK is suppressed)

The WHERE Clause examples below assume that the user is connected to the SAS Scalable Performance Data (SPD) Server LIBNAME *foo* and has executed the following SAS code:

```
data foo.a;
do i=1 to 100;
  do j=1 to 100;
    do k=1 to 100;
      m=mod(i,3);
      output;
    end;
  end;
end;
run;

proc datasets lib=foo;
modify a;
index create ijk = (i j k);
index create j;
index create m;
quit;
```

Example 1 "where i = 1 and j = 2 and m = 4"

```
whinit: WHERE ((I=1) and (J=2) and (M=4))
whinit: wh-tree presented

      /-NAME = [I]
        /-CEQ-----|
        |
      \-LITN = [1]
--LAND---|
        |
      /-NAME = [J]
        |--CEQ-----|
        |
      \-LITN = [2]
        |
      /-NAME = [M]
        \-CEQ-----|

          \-LITN = [4]
whinit: wh-tree after split
--[empty]
whinit: pruning INDEX node which is trivially FALSE
      /-NAME = [M] INDEX M (M)
--CEQ-----|
          \-LITN = [4]
whinit: INDEX tree evaluated to FALSE
whinit returns: FALSE
```

Here the only values that column M can contain are 0, 1, or 2. Thus, the predicate $m = 4$ is identified as trivially FALSE. Because this predicate is part of an AND predicate, it too is FALSE. Consequently, the entire WHERE clause is pre-evaluated to FALSE, meaning that no records can satisfy this WHERE clause. Thus, as a result of the pre-evaluation, no records are actually read from disk. This is an example of optimization at its best.

Example 2: "where i in (1, 2, 3) and j in (4, 5, 6, 7) and k > 8 and m = 2"

```
whinit: WHERE (I in (1, 2, 3) and J in (4, 5, 6, 7) and (K>8) and (M=2))
whinit: wh-tree presented

      /-NAME = [I]
        /-IN-----|
        |
      /-LITN = [1]
        |
      \-SET-----|
        |
          |--LITN = [2]
            |
              \-LITN = [3]
                |
--LAND---|
      /-NAME = [J]
```

```

      |---IN-----|
|      /-LITN = [4]
|      |
\--SET-----|
|      |---LITN = [5]
|      |---LITN = [6]
|      |---LITN = [7]
|      |
/-NAME = [K]
|      |---CGT-----|
|      |
\--LITN = [8]
|      |
/-NAME = [M]
|      |---CEQ-----|

      \-LITN = [2]
whinit: SBM-INDEX M uses 60% of segs(WITHIN maxsegratio 100%)
whinit: wh-tree after split
      /-NAME = [K]
--CGT-----|
      \-LITN = [8]
whinit: INDEX tree after split

      /-NAME = [I] <1>SBM-INDEX IJK (I,J)
      /-IN-----|
|      /-LITN = [1]
|      |
\--SET-----|
|      |---LITN = [2]
|      |---LITN = [3]
--LAND-----|
|      /-NAME = [J]
|      |---IN-----|
|      /-LITN = [4]
|      |
\--SET-----|
|      |---LITN = [5]
|      |---LITN = [6]
|      |---LITN = [7]
|      |
/-NAME = [M] <2>SBM-INDEX M (M)
|      |---CEQ-----|

      \-LITN = [2]
whinit returns: ALL EVAL1(w/SEGLIST) EVAL2

```

Here, a composite index **ijk** was defined on columns (**i j k**). This composite index is used for column's **i** and **j**, which is an equality index predicate. Column **k** is **not** included because it involves a inequality operator (greater than). Since there are no other indexes for column **k**, this predicate is assigned to EVAL2 . EVAL2 will post-filter the records obtained through the use of indexes.

Example 3: "where i = 1 and j > 5 and mod(k, 3) = 2"

```
whinit: WHERE ((I=1) and (J>5) and (MOD(K, 3)=2))
whinit: wh-tree presented
```

```

    /-NAME = [I]
      /-CEQ-----|
      |
    \-LITN = [1]
--LAND---|
      |
    /-NAME = [J]
      |--CGT-----|
      |
    \-LITN = [5]
      |
      /-FUNC = [MOD()]
      |
    /-FLST---|
    |         |
    |         |--NAME = [K]
    |         |
    |         \-LITN = [3]
    |         \-CEQ-----|

```

```

    \-LITN = [2]
whinit: wh-tree after split
```

```

    /-FUNC = [MOD()]
      /-FLST---|
      |
    |--NAME = [K]
      |
    \-LITN = [3]
--CEQ-----|
    \-LITN = [2]

```

```
whinit: SBM-INDEX IJK uses 1% of sges(WITHIN maxsegratio 75%)
whinit: SBM-INDEX J uses at least 76% of segs(EXCEEDS maxsegratio 75%)
whinit: INDEX tree after split
```

```

    /-NAME = [I] <1>SBM-INDEX IJK (I)
      /-CEQ-----|
      |
    \-LITN = [1]
--LAND---|
      |
    /-NAME = [J] <2>SBM-INDEX J (J)
      \-CGT-----|

```

```

    \-LITN = [5]
whinit returns: ALL EVAL1(w/SEGLIST) EVAL2
```

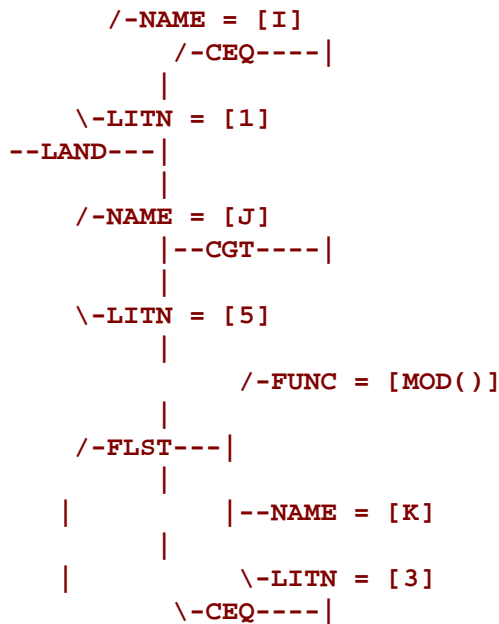
Here the indexes on column **i**, a composite index on the columns (**i j k**), and the column **j** are combined. In this example WHINIT uses both EVAL1 and EVAL2. The **j** predicate involves a inequality operator (greater than). Therefore, WHINIT cannot combine the predicate with **i** and the composite index involving **i** and **j** (and **k**).

Using the composite index **ijk** in this plan may be inefficient. If a smaller composite index (that is, one on **ij** or a simple index on **i**) were available, WHINIT would select it. In lieu of this, try benchmarking the plan. Suppress the composite index and compare the results to the existing plan to see which is more efficient (faster) on your machine.


The example that follows shows what WHINIT's plan would look like with the composite index suppressed.

Example 4: "where i = 1 and j > 5 and mod(k, 3) = 2" (the index IJK is suppressed)

```
whinit: WHERE ((I=1) and (J>5) and (MOD(K, 3)=2))
whinit: wh-tree presented
```



```
\-LITN = [2]
whinit: wh-tree after split
```



```

        \-LITN = [2]
whinit: SBM_INDEX J uses at least 76% of segs (EXCEEDS maxsegratio 75%)
whinit: checking all hybrid segments
whinit: INDEX tree after split
        /-NAME = [J] <1>SBM-INDEX J (J)
--CGT----|
        \-LITN = [5]
whinit returns: ALL EVAL1 EVAL2

```

Notice that the predicate involving column *i* is non-indexed. WHINIT evaluates it using EVAL2. Because the predicate *j* > 5 still uses an inequality comparison, WHINIT continues to use EVAL1. Finally, because the percentage of segments that contain values for column *J* exceeds the maximum segment ratio, pre-segment logic is not done on column *J*. As a result, all segments of the table are queried for values which satisfy the WHERE Clause for column *J*.

Server-Side Sorting

In most instances, using a BY-clause in SAS code submitted to a SAS Scalable Performance Data (SPD) Server table triggers a BY-clause evaluation by SAS Scalable Performance Data (SPD) Server. This BY-clause assertion to the SAS Scalable Performance Data (SPD) Server may or may not require sorting to produce the ordered row set that the BY-clause requires. In some cases, a table index can be used to sort the rows to satisfy a BY-clause.

For example, the input table to a PROC SORT step is sorted in server context (by the associated LIBNAME proxy). The rows are returned to PROC SORT in BY-clause order. In this case, PROC SORT knows that the data is already ordered, and writes the data to the output table without sorting it again. Unfortunately, this approach still must send the data from the LIBNAME proxy to the SAS client and then back to the LIBNAME proxy. However, there are other ways to use a SAS Scalable Performance Data (SPD) Server SQL pass-through COPY statement to avoid the overhead of the data round-trip.

SAS Scalable Performance Data (SPD) Server attempts to use an index when performing a BY-clause. The software looks specifically for an index that has variables in the order specified in the BY-clause. On the surface this seems like a good idea: table row order is already determined because the keys in the index are ordered. SAS Scalable Performance Data (SPD) Server reads the keys in order from the index, then returns the rows from the table based upon the row IDs that are stored with the index key values.

Use caution when using BY-clauses on tables that have indexes on their BY columns. Using the index is not always a good idea. When no suitable index exists to determine BY-clause order, SAS Scalable Performance Data (SPD) Server uses a parallel table scan sort that keeps the table row intact with the sort key. The time required to access a highly random distribution of row IDs (obtained by using the index) can greatly exceed the time required to sort the rows from scratch.

When you use a WHERE Clause to filter the rows from a SAS Scalable Performance Data (SPD) Server table with a BY-clause to order them in a desired way, SAS Scalable Performance Data (SPD) Server handles both the subsetting and the ordering for this request. In this case, the filtered rows that were qualified by the WHERE Clause are fed directly into a sort step. Feeding the filtered rows into the sort step is part of the parallel WHERE Clause evaluation. The final ordered row set is the result. In this case, the previous discussion of index use does not apply. Index use for WHERE Clause filtering is very desirable and greatly improves the filtering performance that feeds into the sort step. Arbitrarily suppressing index use with a WHERE and BY combination should be avoided.

- [Suppressing the Use of Indexes](#)
- [Advantages of Implicit Server Sorts](#)

Suppressing the Use of Indexes

Suppress the use of indexes on the BY-clause by using the SPDSNIDX=YES macro variable or by asserting the NOINDEX=YES table option. Suppressing the use of the index may significantly improve time required to process a BY-clause in SAS Scalable Performance Data (SPD) Server.

Advantages of Implicit Server Sorts

An exceptional feature is the software's ability to execute ad-hoc order-BY queries without pre-sorting the table on the BY variables. Many SAS job streams are structured with code that alternates PROC SORT followed by PROC xxxx invocations where the PROC SORT step is only needed for the execution of the PROC xxxx step.

When sort order is relevant only to the following step, eliminate the PROC SORT step and just use the BY-clause on the PROC xxxx step. This eliminates the extra data transfer (to PROC SORT from SAS Scalable Performance Data (SPD) Server and then back from PROC SORT to SAS Scalable Performance Data (SPD) Server) to store the sorted result. Even if SAS Scalable Performance Data (SPD) Server performs the sort associated with the PROC SORT, there is extra data transfer. The data's round trip from the server to the SAS client and back can impose a substantial time penalty.

Copyright © 2008 by SAS Institute Inc., Cary, NC, USA. All rights reserved.

SAS Scalable Performance Data (SPD) Server Macro Variables

- [Introduction](#)
- [Variable for Compatibility with the Base SAS Engine](#)
 - [SPDSBNEQ=](#)
- [Variables for Miscellaneous Functions](#)
 - [SPDSEOBS=](#)
 - [SPDSSOBS=](#)
 - [SPDSUSAV=](#)
 - [SPDSUSDS=](#)
 - [SPDSVERB=](#)
 - [SPDSFSAV=](#)
 - [SPDSEINT=](#)
- [Variables for Sorts](#)
 - [SPDSBSRT=](#)
 - [SPDSNBIX=](#)
 - [SPDSSTAG=](#)
- [Variables for WHERE Clause Evaluations](#)
 - [SPDSTCNT=](#)
 - [SPDSEV1T=](#)
 - [SPDSEV2T=](#)
 - [SPDSWDEB=](#)
 - [SPDSIRAT=](#)
 - [SPDSNIDX=](#)
 - [SPDSWCST=](#)
 - [SPDSWSEQ=](#)
- [Variables That Affect Disk Space](#)
 - [SPDSCMPF=](#)
 - [SPSDCMP=](#)
 - [SPDSIASY=](#)
 - [SPDSSIZE=](#)
- [Variables To Enhance Performance](#)
 - [SPDSNETP=](#)
 - [SPDSSADD=](#)
 - [SPDSSYRD=](#)
 - [SPDSAUNQ=](#)
- [Variables for a Client and a Server Running on the Same UNIX Machine](#)
 - [SPDSCOMP=](#)

[Introduction](#)

Macro variables, known as symbolic variables, operate similarly to LIBNAME and table options. But, they have an advantage because they apply globally. That is, their value remains constant until explicitly changed.

This chapter presents reference information for SPD Server macro variables, including their purpose, default values, and when and how to use them. The variables are grouped by function or purpose of the default value. Changing the value can also change the purpose, making the variable fall into another group.

For example, the default setting for the macro variable SPDSADD= is NO. The SPDSADD= macro enhances performance during data appends. Setting SPDSADD= to YES changes the way the variable functions. The macro setting SPDSADD=YES ensures compatibility with the base SAS engine. The default setting improves performance. Changing the setting from the default improves Base SAS software compatibility.

To set a macro variable to YES submit the following statement:

```
%let MACROVAR=YES;
```

Note: Assignments for macro variables with YES|NO arguments must be entered in uppercase (capitalized).

When you specify table option settings, precedence matters. If you specify a table option after you set the option in a macro variable statement, the table option setting takes precedence over the macro variable option setting. If you specify an option using a LIBNAME statement, then later specify an option setting through a macro variable statement, the table option setting made in the macro variable takes precedence over the over the LIBNAME statement setting.

To view the default values for the SPD Server macro variables, use the [SPDSMAC](#) command associated with PROC SPDO. SAS displays the macro variables and their current settings. Understanding proper use of macro variables in SPD Server allows you to unleash the power of the software.

Variable for Compatibility with the Base SAS Engine

SPDSBNEQ=

Use the SPDSBNEQ= setting to specify the output order of table rows that have identical values in the BY column.

Syntax

SPDSBNEQ=YES|NO

Default: NO

Corresponding Table Option: BYNOEQUALS=

Use the following arguments:

YES

outputs rows with identical values in a BY clause in random order.

NO

outputs rows with identical values in a BY clause using the relative table position of the rows from the input table.

Description

SPDSBNEQ=NO configures the SPD Server to imitate the Base SAS engine behavior. If strict compatibility is not required, assign SPDSBNEQ=YES. Random output allows the SPD Server to create indexes and append to tables faster.

Example

Configure the SPD Server so that it output table rows as quickly as possible when processing rows that have identical values in the BY column.

```
%let SPDSBNEQ=YES;
```

Variables for Miscellaneous Functions

SPDSEOBS=

Use the SPDSEOBS= macro variable to specify the number of the last row (end observation) of a user-defined range that you want to process in a table.

Syntax

```
SPDSEOBS=n
```

Default: The default setting of 0 processes the entire table.

Corresponding Table Option: ENDOBS=

Use the following argument:

n
is the number of the end row.

Description

The SPD Server processes the entire table by default unless you specify a range of rows. You can specify a range using the macro variables SPDSOBS= and SPDSEOBS=, or you can use the table options, STARTOBS= and ENDOBS=.

If you use the range start macro variable SPDSOBS= without specifying an end range value using the SPDSEOBS= macro variable, SPD Server processes to the last row in the table. If you specify values for both SPDSOBS= and SPDSEOBS= macro variables, the value of SPDSEOBS= must be greater than SPDSOBS=. The SPDSOBS= and SPDSEOBS= macro variables specify ranges for table input processing as well as WHERE Clause processing.

Example

In order to create test tables, you configure the SPD Server to subset the first 100 rows of each table in your job. Submit the macro variable statement for SPDSEOBS= at the beginning of your job.

```
%let SPDSEOBS=100;
```

SPDSOBS=

Use the SPDSOBS= macro variable to specify the number of the starting row (observation) in a user-defined range of a table.

SyntaxSPDSSOBS=*n***Default:** The default setting of 0 processes the entire table.**Corresponding Table Option:** STARTOBS=

Use the following argument:

n

is the number of the start row.

Description

By default, SPD Server processes entire tables unless you specify a range of rows. You can specify a range using the macro variables SPDSSOBS= and SPDSEOBS=, or you can use the table options, STARTOBS= and ENDOBS=.

If you specify the end of a user-defined range using the SPDSEOBS= macro variable, but do not implicitly specify the beginning of the range using SPDSSOBS=, SPD Server sets SPDSSOBS= to 1, or the first row in the table. If you specify values for both SPDSSOBS= and SPDSEOBS= macro variables, the value of SPDSEOBS= must be greater than SPDSSOBS=. The SPDSSOBS= and SPDSEOBS= macro variables specify ranges for table input processing as well as WHERE Clause processing.

Example

Print the INVENTORY.OLDAUTOS table, skipping rows 1-999 and beginning with row 1000. You should submit the SPDSSOBS= macro variable statement before the PROC PRINT statement in your job.

```
%let SPDSSOBS=1000;
```

The statement above specifies the starting row with SPDSSOBS=, but does not declare an ending row for the range using SPDSEOBS=. When the program executes, SAS will begin printing at row 1000 and continues until the final row of the table is reached.

```
PROC PRINT data=inventory.oldautos;  
run;
```

SPDSUSAV=

Use the SPDSUSAV= macro variable to specify whether to save rows with non-unique (rejected) keys to a separate SAS table.

Syntax

SPDSUSAV=YES|NO|REP

Default: NO**Affected by Table Option :** SYNCADD=

Use in Conjunction with Variable : SPDSUSDS=

Corresponding Table Option : UNIQUESAVE=

Use the following arguments:

YES

writes rows with non-unique key values to a SAS table. Use the macro variable SPDSUSDS= to reference the name of the SAS table for the rejected keys.

NO

non-unique key values are ignored and rejected rows are not written to a separate table.

REP

when updating a master table from a transaction table, where the two tables share identical variable structures, the SPDSUSAV=REP option replaces the row updated row in the master table instead of appending a row to the master table. The REP option only functions in the presence of a /UNIQUE index on the MASTER table. Otherwise, the REP setting is ignored..

Description

When performing an append operation, SPD Server does not save the rows which contain duplicate key values unless the SPDSUSAV= macro variable is set to YES.

When SPDSUSAV= is set to YES, SPD Server creates a hidden SAS table and writes rejected rows to the table. Use the SPDSUSDS= macro variable command to view the contents of the table. Each append operation creates a different table.

Example

Append several tables to the EMPLOYEE table, using employee number as a unique key. The appended tables should not have records with duplicate employee numbers.

At the beginning of the job, configure SPD Server to write any rejected (identical) employee number records to a SAS table. The macro variable SPDSUSDS= holds the name of the SAS table for the rejected keys.

```
%let SPDSUSAV=YES
```

Use a %PUT statement to display the name of the table, and then print the table.

```
%put Set the macro variable spdsusds to &spdsusds;  
  
title 'Duplicate (non-unique) employee numbers found in  
EMPS';  
PROC PRINT data=&spdsusds; run;
```

SPDSUSDS=

Use the SPDSUSDS= macro variable to reference the name of the SAS table that SPD Server creates for duplicate or rejected keys when the SPDSUSAV= macro variable is set to YES.

Syntax

SPDSUSDS=

Default: SPD Server automatically generates identifying strings for the duplicate or rejected key tables.

Use in Conjunction with Table Option: SYNCADD=

Use in Conjunction with Variable: SPDSUSAV=

Corresponding Table Option: UNIQUESAVE=

Description

When SPDSUSAV= or UNIQUESAVE= is set to YES, SPD Server creates a table to store any rows with duplicate key values encountered during an append operation. Submitting the SPDSUSDS= macro variable references the generated name for the hidden SAS table.

To obtain the name and print the table's contents, reference the variable SPDSUSDS=.

Example

```
%let SPDSUSAV=YES
```

Use a %PUT statement to display the name of the table created by SPDSUSDS= and to print out the duplicate rows.

```
%put Set the macro variable spdsusds to &spdsusds;  
  
title 'Duplicate Rows Found in MYTABLE  
      During the Last Data Append';  
PROC PRINT data=&spdsusds; run;
```

SPDSVERB=

Use the SPDSVERB= macro variable to provide verbose details on all indexes, ACL information, and other information that is associated with SPD Server tables.

Syntax

SPDSVERB=YES|NO

Default: NO

Corresponding Table Option: VERBOSE=

Use the following arguments:

YES

requests detail information for indexes, ACLs, and other SPD Server table values.

NO

suppresses detail information for indexes, ACLs, and other SPD Server table values.

Example

You need information about associated indexes for the SPD Server table SUPPLY. Configure SPD Server for verbose details at the start of your session so you can see index details. Submit the SPDSVERB= macro variable as a line in your your autoexec.sas file:

```
%let SPDSVERB=YES;
```

Submit a PROC CONTENTS request for the SUPPLY table:

```
PROC CONTENTS data=supply;  
run;
```

SPDSFSAV=

Use the SPDSFSAV= macro variable to specify whether you want to retain table data if the SPD Server table creation process terminates abnormally.

Syntax

SPDSFSAV=YES|NO

Default: NO. Normally SAS closes and deletes tables which are not properly created.

Use the following arguments:

YES

enables FORCESAVE mode and saves the table.

NO

default SPD Server actions delete partially completed tables.

Description

Large tables can require a long time to create. If problems such as network interruptions or disk space shortages occur during this time period, the table may not be properly created and signal an error condition. If SAS encounters such an error condition, it deletes the partially completed table.

In SPD Server you can set SPDSFSAV=YES. Saving the partially created table can protect the time and resources invested in a long-running job. When the SPDSFSAV= macro variable is set to YES, the SPD Server LIBNAME proxy saves partially completed tables in their last state and identifies them as damaged tables.

Marking the table damaged prohibits other SAS DATA or PROC steps from accessing the table until its state of completion can be verified. After you verify or repair a table, you can clear the 'damaged' status and enable further read/update/append operations on the table. Use the PROC DATASETS REPAIR operation to remove the damaged file indicator.

Example

Configure SPD Server before you run the table creation job for a large table called ANNUAL. If some error

prevents the successful completion of the table ANNUAL, the partially completed table will be saved.

```
%let SPDSFSAV=YES;  
DATA SPDSLIB.ANNUAL;  
...  
RUN;
```

SPDSEINT=

Use the SPDSEINT= macro to specify how SPD Server responds to network disconnects during SQL pass-through EXECUTE() statements.

Syntax

SPDSEINT=YES|NO

Default: YES

Description:

The SPD Server SQL server interrupts SQL processing by default when a network failure occurs . The interruption prematurely terminates the EXECUTE() statement. Setting SPDSEINT=NO configures the SPD Server's SQL server to continue processing until completion regardless of network disconnects.

Warning: Use the macro variable setting SPDSEINT=NO carefully! A runaway EXECUTE() statement requires a privileged system user on the server machine to kill the SPD Server SQL proxy process. This is the only way to stop the processing.

Variables for Sorts

SPDSBSRT=

Use the SPDSBSRT= macro variable to configure SPD Server's sorting behavior when it encounters a BY-clause and there is no index available.

Syntax

SPDSBSRT=YES|NO

Default: YES

Corresponding Table Option: BYSORT=

Use the following arguments:

YES

SPD Server performs a server sort when it encounters a BY clause and there is no index available.

NO

SPD Server does not perform a sort when it encounters a BY clause.

Description

Base SAS software requires an explicit PROC SORT statement to sort SAS data. In contrast, SPD Server sorts a table whenever it encounters a BY clause, if it determines that the table has no index.

Advantages for using SPD Server implicit sorts are discussed in detail in the Help section for [Connecting to SAS Scalable Performance Data Server](#).

Example 1

At the start of a session to run old SAS programs, you realize that you do not have time to remove the existing PROC SORT statements. These statements are present only to generate print output.

To avoid redundant Server sorts, configure SPD Server to turn off implicit sorts. Put the macro variable assignment in your autoexec.sas file so SPD Server retains the configuration for all job sessions.

```
%let SPDSBSRT=NO;
```

During the Example 1 session you decide to run a new program that has no PROC SORT statements. Instead, the new program takes advantage of SPD Server implicit sorts.

```
data inventory.old_autos;  
  input  
    year $4.  
    @6 manufacturer $12.  
    model $10.  
    body_style $5.  
    engine_liters  
    @39 transmission_type $1.  
    @41 exterior_color $10.  
    options $10.  
    mileage condition;  
  
  datalines;  
  
  1971 Buick      Skylark  conv  5.8  A  yellow  00000001 143000 2  
  1982 Ford      Fiesta   hatch  1.2  M  silver  00000001  70000 3  
  1975 Lancia    Beta     2door  1.8  M  dk blue  00000010  80000 4  
  1966 Oldsmobile Toronado 2door  7.0  A  black   11000010 110000 3  
  1969 Ford      Mustang  sptrf  7.1  M  red     00000111 125000 3  
  ;  
  
  PROC PRINT data=inventory.old_autos;  
    by model;  
  
  run;
```

When the code executes, the PRINT procedure returns an error message. What happened? SAS expected INVENTORY.OLDAUTOS to be sorted before it would generate print output. Since there is no PROC SORT statement -- and implicit sorts are still turned off -- the sort does not occur.

Example 2

Keep implicit sorts turned off for the session, but specify an implicit sort for the table INVENTORY.OLDAUTOS.

```
PROC PRINT data=inventory.oldautos(bysort=yes);  
by model;  
run;
```

SPDSNBIX=

Use the SPDSNBIX= macro variable to configure whether to use an index during a BY-sort.

Syntax

SPDSNBIX=YES|NO

Default: NO

Corresponding Server Parameter Option: [NO]BYINDEX

Use the following arguments:

YES

Set SPDSNBIX=YES to suppress index use during a BY-sort. If the distribution of the values in the table are not relatively sorted or clustered, using the index for the by sort can result in poor performance.

NO

Set SPDSNBIX=NO or use the default value to allow the [NO]BYINDEX server parameter option to determine whether to use an index for a by sort.

Example

```
%let SPDSNBIX=YES;
```

SPDSSTAG=

Use the SPDSSTAG= macro variable to specify whether to use non-tagged or tagged sorting for PROC SORT or BY processing.

Syntax

SPDSSTAG=YES|NO

Default: NO

Use the following arguments:

YES

performs tagged sorting.

NO

performs non-tagged sorting.

Description

During a non-tagged sort, SPD Server attaches the entire table column to the key field(s) to be sorted. Non-tagged sorting allows the software to deliver better performance than a tagged sort. Non-tagged sorting also requires more temporary disk space than a tagged sort.

Example

You are running low on disk space and don't know if you have enough disk overhead to accommodate the extra sort space required to support a non-tagged sort operation.

Configure SPD Server to perform a tagged sort.

```
%let SPDSSTAG=YES;
```

Variables for WHERE Clause Evaluations

SPDSTCNT=

Use the SPDSTCNT= macro variable to specify the number of threads that you want to use during WHERE Clause evaluations.

Syntax

SPDSTCNT=*n*

Default: The value of MAXWHTHREADS is configured by SPD Server parameters.

Used in Conjunction with the SPD Server Parameter: MAXWHTHREADS

Corresponding Table Option: THREADNUM=

Use the following argument:

n

is the number of threads.

Description

See [THREADNUM=](#) for a description and an explanation of how SPDSTCNT= interacts with the SPD Server parameter MAXWHTHREADS.

SPDSEV1T=

Use the SPDSEV1T= macro variable to indicate whether data returned from a SPD Server WHERE Clause evaluations should be in strict row (observation) order.

The macro variables SPDSEV1T= and SPDSEV2T= work in conjunction with the SPD Server WHERE Clause planner WHINIT.

The variables SPDSEV1T= and SPDSEV2T= are identical in purpose. You use them to specify the row order of data returned in WHERE-processing. Which variable the server exercises depends on the evaluation strategy selected by WHINIT. The SPDSEV1T= evaluation strategy is indexed. The SPDSEV2T= evaluation strategy is non-indexed. Avoid using these options unless you absolutely understand the SPD Server performance tradeoffs that depend on maintaining the order of data.

If compatibility with Base SAS software is important, set both SPDSEV1T= and SPDSEV2T= to 0. When both evaluation strategies are set to 0, SPD Server returns data in row order whether the SPDSEV1T= or the SPDSEV2T= strategy is selected.

When you use PROC SQL to perform table joins on sorted tables that use extra WHERE predicates to filter table rows, PROC SQL uses sort order information to optimize the join strategy. Data that has been filtered by the evaluation of extra WHERE predicates returns to PROC SQL in sorted order. SPD Server restricts parallel evaluation of WHERE predicates any time the table(s) have been sorted using PROC SORT to meet this requirement. Sorting data in this manner can negatively impact net performance because most SAS PROCs or DATA steps do not process sorted data any faster than unsorted data.. SPD Server recommends you set the SPDSEV1T= and/or SPDSEV2T= macro variables to 2. Setting SPDSEV1T= and/or SPDSEV2T= macro variables to 2 configures SPD Server to perform parallel WHERE Clause evaluations without regard for the sort order of the SPD Server tables.

Note: The SPDSEV1T= and SPDSEV2T= usage discussed here does not apply to SQL statements executed via the SPD Server pass-through SQL mechanism.

Syntax

SPDSEV1T=0|1|2

Default: 1

Used in Conjunction with Indexed WHERE Clause Evaluation Strategy

Uses the following arguments:

- 0** returns data in row order.
- 1** may not return the data in row order. SPD Server may override as needed to force a **0** setting if the table is sorted using PROC SORT.
- 2** always forces parallel evaluation regardless of sorted order. May not return data in row order.

Description

If SPD Server must return many rows during WHERE Clause processing, setting the variable to **0** will greatly slow performance. Use **0** only when row order is required. Use **2** only when you know row order is not important to the result.

Example

Configure SPD Server to send back data in row order whenever WHINIT performs an EVAL1 evaluation.

```
%let SPDSEV1T=0;
```

SPDSEV2T=

Use the SPDSEV2T= macro variable to specify whether the data returned from WHERE Clause evaluations should be in strict row (observation) order.

The macro variables SPDSEV1T= and SPDSEV2T= work in conjunction with the SPD Server WHERE Clause planner WHINIT.

The variables SPDSEV1T= and SPDSEV2T= are identical in purpose. You use them to specify the row order of data returned in WHERE-processing. Which variable the server exercises depends on the evaluation strategy selected by WHINIT. The SPDSEV1T= evaluation strategy is indexed. The SPDSEV2T= evaluation strategy is non-indexed. Avoid using these options unless you absolutely understand the SPD Server performance tradeoffs that depend on maintaining the order of data.

If compatibility with Base SAS software is important, set both SPDSEV1T= and SPDSEV2T= to 0. When both evaluation strategies are set to 0, SPD Server returns data in row order whether the SPDSEV1T= or the SPDSEV2T= strategy is selected.

When you use PROC SQL to perform table joins on sorted tables that use extra WHERE predicates to filter table rows, PROC SQL uses sort order information to optimize the join strategy. Data that has been filtered by the evaluation of extra WHERE predicates returns to PROC SQL in sorted order. SPD Server restricts parallel evaluation of WHERE predicates any time the table(s) have been sorted using PROC SORT to meet this requirement. Sorting data in this manner can negatively impact net performance because most SAS PROCs or DATA steps do not process sorted data any faster than unsorted data.. SPD Server recommends you set the SPDSEV1T= and/or SPDSEV2T= macro variables to 2. Setting SPDSEV1T= and/or SPDSEV2T= macro variables to 2 configures SPD Server to perform parallel WHERE Clause evaluations without regard for the sort order of the SPD Server tables.

Note: The SPDSEV1T= and SPDSEV2T= usage discussed here does not apply to SQL statements executed via the SPD Server pass-through SQL mechanism.

Syntax

SPDSEV2T=0|1|2

Default: 1

Used in Conjunction with Non-Indexed WHERE Clause Evaluation Strategy

Use the following arguments:

- 0** returns data in row order.
- 1** may not return the data in row order. SPD Server may override as needed to force **0** setting if the table is sorted using PROC SORT.
- 2** always forces parallel evaluation regardless of sorted order. May not return the data in row

order.

Description

If SPD Server must return many rows during WHERE Clause processing, setting the variable to **0** will greatly slow performance. Use **0** only when row order is required. Use **2** only when you know row order is not important to the result.

Example

Configure SPD Server to send back data in row order whenever WHINIT performs an EVAL2 evaluation.

```
%let SPDSEV2T=0;
```

SPDSWDEB=

Use the SPDSWDEB= macro variable to specify whether the WHERE Clause planner WHINIT, when evaluating a WHERE expression, should display a summary of the execution plan.

Syntax

SPDSWDEB=YES|NO

Default: NO

Use the following arguments:

YES

displays WHINIT's planning output.

NO

suppresses WHINIT's planning output.

SPDSIRAT=

Use the SPDSIRAT= macro variables to specify whether to perform segment candidate pre-evaluation when performing WHERE Clause processing with hybrid indexes.

Syntax

SPDSIRAT=0..100

Default: MAXSEGRATIO server parameter

Description:

When using hybrid indexes, WHERE-based queries pre-evaluate segments. The segments are scanned for candidates that match one or more predicates in the WHERE clause. The candidate segments that were identified during the pre-evaluation are queried in subsequent logic to evaluate the WHERE Clause. Eliminating the non-candidate segments from the WHERE Clause evaluation

generally results in substantial performance gains.

Some queries can benefit by limiting the pre-evaluation phase. SPD Server imposes the limit based on a ratio: the number of segments that contain candidates compared to the total number of segments in the table. The reason for this is simple. If the predicate has candidates in a high percentage of the segments, the pre-evaluation work is largely wasted.

The ratio formed by dividing the number of segments that containing candidates by the number of total segments is compared to a cutoff point. If the segment ratio is greater than the value assigned to the cutoff point, the extra processing required to perform pre-evaluation outweighs any potential process savings that might be gained through the predicate pre-evaluation. SPD Server calculates the ratio for a given predicate and compares the ratio to the SPDSIRAT= value, which acts as the cutoff point. If the calculated ratio is less than or equal to the SPDSIRAT= value, pre-evaluation is performed. If the calculated ratio is greater than the SPDSIRAT= value, pre-evaluation is skipped and every segment is a candidate for the WHERE Clause.

Use the global SPD Server parameter, MAXSEGRATIO to set the default cutoff value. The default MAXSEGRATIO should provide good performance. Certain specific query situations might be justification for modifying your SPDSIRAT= value. When you modify your SPDSIRAT= value, it overrides the default value established by MAXSEGRATIO.

Example:

Configure SPD Server to perform a pre-evaluation phase for WHERE Clause processing with hybrid indexes if the candidates are in 65% or less of the segments.

```
%let SPDSIRAT=65;
```

SPDSNIDX=

Use the SPDSNIDX= macro variable to specify whether to use the table's indexes when processing WHERE Clauses. SPDSNIDX= can also be used to disable index use for BY- order determination.

Syntax

SPDSNIDX=YES|NO

Default: NO

Corresponding Table Option: NOINDEX=

Use the following arguments:

YES

ignores indexes when processing WHERE Clauses.

NO

uses indexes when processing WHERE Clauses.

Description:

Set SPDSNIDX=YES to test the effect of indexes on performance or for specific processing. Do not

use YES routinely for normal processing.

Example:

Assume you are processing data from SPORT.MAILLIST. There is an index for the SEX column and you want to test to determine if the index will improve performance when you use PROC PRINT processing on SPORT.MAILLIST.

You should configure SPD Server not to use the index:

```
data sport.maillist;
  input
    name $ 1-20
    address $ 21-57
    phoneno $ 58-69
    sex $71;

datalines;

  Douglas, Mike 3256 Main St., Cary, NC 27511 919-444-5555 M
  Walters, Ann Marie 256 Evans Dr., Durham, NC 27707 919-324-
6786 F
  Turner, Julia 709 Cedar Rd., Cary, NC 27513 919-555-9045 F
  Cashwell, Jack 567 Scott Ln., Chapel Hill, NC 27514 919-533-
3845 M
  Clark, John 9 Church St., Durham, NC 27705 919-324-0390 M
  ;

PROC DATASETS lib=sport nolist;
modify maillist;
index create sex;
quit;

/*Turn on the macro variable SPDSWDEB */
/* to show that the index is not used */
/* during the table processing.      */

%let spdswdeb=YES;

%let spdsnidx=YES;

title "All Females from Current Mailing List";
PROC PRINT data=sport.maillist;
where sex="F";
run;

%let spdsnidx=NO;
```

SPDSWCST=

Use the SPDSWCST= macro variable to specify whether to use dynamic WHERE Clause costing.

Syntax

SPDSWCST=YES|NO

Default: NO

Corresponding Server Parameter Option: [NO]WHERECOSTING

Turns WHERE-costing on or off for an entire server.

Description:

Set SPDSWCST=YES to use dynamic WHERE Clause costing. Disabling SPDSWCST= defaults SPD Server to using WHERE-costing with WHINIT.

Example:

```
%let SPDSWCST=YES;
```

SPDSWSEQ=

Syntax

SPDSWSEQ=YES|NO

Default: NO

Description:

Set the SPDSWSEQ= macro variable to YES. When set to YES, the SPDSWSEQ= macro variable overrides WHERE Clause costing and forces a global EVAL3 or EVAL4 strategy.

Example:

```
%let SPDSWSEQ=YES;
```

Variables That Affect Disk Space

SPDSCMPF=

Use the SPDSCMPF= macro variable to specify the amount of growth space, sized in bytes, to be added to a compressed data block.

Syntax

SPDSCMPF=*n*

Default: 0 bytes

Use the following argument:

n

is the number of bytes to add.

Description

Updating rows in compressed tables can increase the size of a given table block. Additional space is required for the block to be written back to disk. When contiguous space is not available on the hard drive, a new block fragment stores the excess, updated quantity. Over time, the table will experience block fragmentation.

When opening compressed tables for OUTPUT or UPDATE, you can use the SPDSCMPF= macro variable to anticipate growth space for the table blocks. If you estimate correctly, you can greatly reduce block fragmentation in the table.

Note: SPD Server table metadata does not retain compression buffer or growth space settings.

SPDSDCMP=

Use the SPDSDCMP= macro variable to compress SPD Server tables that are stored on disk.

Syntax

SPDSDCMP=YES|NO

Default: NO

Use in Conjunction with Table Option: IOBLOCKSIZE=

Corresponding Table Option: COMPRESS=

Use the following arguments:

YES

performs the run-length compression algorithm SPDSRLLC.

NO

performs no table compression.

Description

When you set the SPDSDCMP= macro variable to YES, SPD Server compresses newly created tables by 'blocks' according to the algorithm specified. To control the amount of compression, use the table option IOBLOCKSIZE= to specify the number of rows that you want to store in the block. For a complete discussion, refer to [IOBLOCKSIZE=](#).

Note: Once a compressed table is created, you cannot change its block size. To resize the block, you must PROC COPY the table to a new table, setting IOBLOCKSIZE= to the new block size for the output table.

Example

Prior to creating huge tables, you want to conserve disk space. Specify compression, and the default algorithm SPDSRLLC, at the beginning of your job.

```
%let SPDSDCMP=YES;
```

SPDSIASY=

Use the SPDSIASY= macro variable to specify whether to create indexes in parallel when creating multiple indexes on an SPD Server table.

Syntax

SPDSIASY=YES|NO

Default: NO

Corresponding Table Option : ASYNCINDEX=

Use the following arguments:

YES

creates the indexes in parallel.

NO

creates one index at a time.

Description

You use the macro variable SPDSIASY= to choose between parallel and sequential index creation on SPD Server tables with more than one index. One advantage of creating multiple indexes in parallel is speed. The speed enhancements that can be achieved with parallel indexes are not free. Parallel indexes require significantly more disk space for working storage. The default SPD Server setting for the SPDSIASY= macro variable is set to NO, in order to avoid exhausting the available work storage space.

However, if you have adequate disk space to support parallel sorts, **it is strongly recommended** that you override the default SPDSIASY=NO setting and assign SPDSIASY=YES. You can substantially increase performance -- indexes that take hours to build complete much faster.

How many indexes should you create in parallel? The answer depends on several factors, such as the number of CPUs in the SMP configuration and available storage space needed for index key sorting.

When managing disk space on your SPD Server, remember that grouping `index create` statements can minimize the number of table scans that SPD Server performs, but it also affects disk space consumption. There is an inverse relationship between the table scan frequency and disk space requirements. A minimal number of table scans requires more auxiliary disk space; a maximum number of table scans requires less auxiliary disk space.

Example

Your perform batch processing from midnight to 6:00 a.m. All of your processing must be completed before start of the next work day. One frequently-repeated batch job creates large indexes on a table, and usually takes several hours to complete. Configure SPD Server to create indexes in parallel to reduce the processing time.

```
%let SPDSIASY=YES;
proc datasets lib=spds;
  modify a;
  index create x;
  index create y;
  modify a;
  index create comp=(x y) comp2=(y x);
quit;
```

In the example above, the X and Y indexes will be created in parallel. After creating X and Y indexes, SPD Server creates the COMP and COMP2 indexes in parallel. In this example, two table scans are required: one table scan for the X and Y indexes, and a second table scan for the COMP and COMP2 indexes.

SPDSSIZE=

Use the SPDSSIZE= macro variable to specify the size of an SPD Server table partition.

Syntax

SPDSSIZE=*n*

Default: 16 Megabytes

Corresponding Table Option: PARTSIZE=

Affected by LIBNAME option: DATAPATH=

Use the following argument:

n

is the size of the partition in Megabytes.

Description

Use this SPDSSIZE= macro variable option to improve performance of WHERE Clause evaluation on non-indexed table columns.

Splitting the data portion of a server table at fixed-sized intervals allows SPD Server to introduce a high degree of scalability for non-indexed WHERE Clause evaluation. This is because SPD Server launches threads in parallel and can evaluate different partitions of the table without file access or thread contention. The speed enhancement comes at the cost of disk usage. The more data table splits you create, the more you increase the number of files, which are required to store the rows of the table.

Scalability limits on the SPDSSIZE= macro variable ultimately depend on how you structure the DATAPATH= option in your LIBNAME statement. The configuration of the DATAPATH= file systems across striped volumes is important. You should spread each individual volume's striping configuration across multiple disk controllers/SCSI channels in the disk storage array. Your configuration goal, at the hardware level, should be to maximize parallelism when performing data retrieval.

The SPDSSIZE= specification is also limited by MINPARTSIZE=, an SPD Server parameter maintained by the SPD Server administrator. MINPARTSIZE= ensures that an over-zealous SAS user cannot arbitrarily create small partitions, thereby generating an excessive number of physical files. The default for MINPARTSIZE= is 16 Mbytes.

Note: The SPDSSIZE= value for a table cannot be changed after the table is created. To change the SPDSSIZE=, you must PROC COPY the table and use a different SPDSSIZE= (or PARTSIZE=) option setting on the new (output) table.

For an example using the table option, see [PARTSIZE=](#).

```
%let SPDSSIZE=32;
```

Variables To Enhance Performance

SPDSNETP=

Use the SPDSNETP= macro variable to size buffers in server memory for the network data packet.

Syntax

SPDSNETP=*size-of-packet*

Default: 32K

Corresponding Table Option: [NETPACKSIZE=](#)

Use the argument:

size-of-packet

is the size (integer) in bytes of the network packet.

Description

When sizing the buffer for data packet transfer between SPD Server and your SAS client machine, the packet must be greater than or equal in size to one table row. See [NETPACKSIZE=](#) for more information.

Example

Despite recent upgrades to your network connections, you are experiencing significant pauses when the SPD Server transfers data. You want to resize the data packet to send three rows at a time for a more continuous data flow.

Specify a buffer size in server memory that is three times the row size (6144 bytes.) Submit your SPDSNETP= macro variable statement at the top of your job.

```
%let SPDSNETP=18432;
```

SPDSSADD=

Use the SPDSSADD= macro variable to specify whether SPD Server appends tables by transferring a single row at a time synchronously, or by transferring multiple rows asynchronously (block row appends).

Syntax

SPDSSADD=YES|NO

Default: NO

Related Table Option: SYNCADD=

Use the following arguments:

YES

applies a single row at a time during an append operation. This behavior imitates the Base SAS engine.

NO

appends multiple rows at a time

Description

SPDSSADD=YES slows performance. Use this argument only if you require strict compatibility with Base SAS software when processing a table. For a complete discussion, refer to [SYNCADD=](#).

SPDSSYRD=

Use the SPDSSYRD= macro variable to specify whether SPD Server should perform asynchronous data streaming when reading a table.

Syntax

SPDSSYRD=YES|NO

Default: NO

Related Table Option: SYNCREAD=

Use the following arguments:

YES

enables asynchronous data streaming.

NO

disables asynchronous data streaming.

Description Use SPDSSYRD=YES only with a MODIFY statement. If you use it with any other processing operation, you slow performance.

SPDSAUNQ=

Use the SPDSAUNQ= macro variable setting to specify whether to abort an append to a table if the table has a unique index and the append would violate the index uniqueness.

Syntax

SPDSAUNQ=YES|NO

Default: NO

Description: Use SPDSAUNQ=YES macro variable to improve append performance to a table with unique indexes. If uniqueness is not maintained, the append is aborted and the table is returned to its state prior to the append. In such an instance, you can scrub the table to remove non-unique values and re-do the append with the macro variable SPDSAUNQ= set to YES. The other alternative is to simply re-do the append with the macro variable SPDSAUNQ= set to NO.

If SPDSAUNQ=NO, the SPD Server will enforce uniqueness at the expense of appending unique indexes in observation order one row at a time. If uniqueness is not maintained for any given row, that row is discarded from the append.

Variables for a Client and a Server Running on the Same UNIX Machine

SPDSCOMP=

specifies to compress the data when sending a data packet through the network.

Syntax

SPDSCOMP=YES|NO

Default: NO

Copyright © 2008 by SAS Institute Inc., Cary, NC, USA. All rights reserved.

SAS Scalable Performance Data (SPD) Server LIBNAME Options

- [Introduction](#)
- [Options to Locate a SAS Scalable Performance Data \(SPD\) Server Host](#)
 - [HOST=](#)
- [Using a Macro Variable to Specify the SAS Scalable Performance Data \(SPD\) Server Host](#)
 - [SERVER=](#)
- [Options to Identify the SAS Scalable Performance Data \(SPD\) Server Client](#)
 - [ACLGRP=](#)
 - [CHNGPASS=](#)
 - [NEWPASSWORD=](#) or [NEWPASSWD=](#)
 - [PASSWORD=](#) or [PASSWD=](#)
 - [PROMPT=](#)
 - [USER=](#)
- [Options to Specify Implicit SQL Pass-Through](#)
 - [IP=YES](#)
 - [PASSTHRU=](#)
- [Options to Specify File Paths for Table Storage](#)
 - [CREATE=](#)
 - [DATAPATH=](#)
 - [INDEXPATH=](#)
 - [METAPATH=](#)
- [Options for Access Control Lists \(ACLs\)](#)
 - [ACLSPECIAL=](#)
- [Options for a Client and Server Running on the Same UNIX Machine](#)
 - [NETCOMP=](#)
 - [UNIXDOMAIN=](#)
- [Options for Other Functions](#)
 - [BYSORT=](#)
 - [DISCONNECT=](#)
 - [ENDOBS=](#)
 - [LIBGEN=](#)
 - [LOCKING=](#)
 - [STARTOBS=](#)
 - [TEMP=](#)
 - [TRUNCWARN=](#)
 - [WORKPATH=](#)

[Introduction](#)

All SAS users who want to use LIBNAME access to a SAS Scalable Performance Data (SPD) Server should read this chapter.

This chapter contains reference information for the SAS Scalable Performance Data (SPD) Server LIBNAME options. The options are grouped by the function or purpose of their default value. You can change the default, thereby controlling how they

function in different data situations. The examples for using the options assume that a LIBNAME statement to access the SAS Scalable Performance Data (SPD) Server engine SASSPDS has previously been issued.

When using the options, remember that if a table option is used subsequent to a LIBNAME option of the same name, the value of the table option or macro variable takes precedence.

Options to Locate a SAS Scalable Performance Data (SPD) Server Host

HOST=

Specifies a SAS Scalable Performance Data (SPD) Server machine by node name or IP address and locates the name server using the SERVICE value.

Syntax

```
HOST='hostname' <SERVICE='service'>
```

Use the following arguments:

'hostname'

is the node name of the SAS Scalable Performance Data (SPD) Server machine or an IP address.

'service'

is the name of a service or the port number for the SAS Scalable Performance Data (SPD) Server's name server.

Description

This option provides the node name of a SAS Scalable Performance Data (SPD) Server host machine and locates the port number of the SAS Scalable Performance Data (SPD) Server's name server. When there is no SERVICE= specification, SAS Scalable Performance Data (SPD) Server checks the client's **/etc/services** file (or its equivalent file) for SPDSNAME -- a reserved name for the SAS Scalable Performance Data (SPD) Server's name server.

Examples

Specify the server machine SAMSON and use the default named service SPDSNAME to obtain the port number of the SAS Scalable Performance Data (SPD) Server's name server.

```
LIBNAME mylib sasspds 'spdsdata'  
host='samson';
```

Specify the server machine SAMSON and provide the port number of the SAS Scalable Performance Data (SPD) Server's name server.

```
LIBNAME mylib sasspds 'spdsdata'  
host='samson'  
service='5002';
```

Using a Macro Variable to Specify the SAS Scalable Performance Data (SPD) Server Host

Assign the macro variable SPDSHOST to the SAS Scalable Performance Data (SPD) Server host SAMSON so that the LIBNAME statement need not specify SAMSON.

```
%let spdshost=samson;
LIBNAME mylib sasspds 'spdsdata'
    user='yourid'
    password='swami';
```

SERVER=

Specifies a SAS Scalable Performance Data (SPD) Server host machine by node name and locates the network address (port number) of the SAS Scalable Performance Data (SPD) Server name server.

Syntax

```
SERVER=hostname.servname
```

Use the following arguments:

hostname

is the node name of the SAS Scalable Performance Data (SPD) Server host machine.

servname

is the name of a service or the port number of the SAS Scalable Performance Data (SPD) Server's name server.

Examples

Specify the SAS Scalable Performance Data (SPD) Server host machine SAMSON and use the default named service SPDSNAME to obtain the port number of the SAS Scalable Performance Data (SPD) Server's name server.

```
LIBNAME mylib sasspds 'spdsdata'
    server=samson.spdsname;
```

Specify the SAS Scalable Performance Data (SPD) Server host machine SAMSON and give the port address of the SAS Scalable Performance Data (SPD) Server's name server.

```
LIBNAME mylib sasspds 'spdsdata'
    server=samson.5002;
```

Options to Identify the SAS Scalable Performance Data (SPD) Server Client

ACLGRP=

Names an ACL group which has been previously assigned to the SAS Scalable Performance Data (SPD) Server user ID. The SAS Scalable Performance Data (SPD) Server system administrator sets up ACL groups and can assign a single user to up to five ACL groups.

Syntax

```
ACLGRP= 'aclgroup'
```

Use the following argument:

aclgroup

Names the ACL group that the SAS Scalable Performance Data (SPD) Server Administrator assigned to your SAS Scalable Performance Data (SPD) Server user ID. (You can be assigned up to five ACL groups.)

Example

Specify the ACL group PROD.

```
LIBNAME mylib sasspds 'spdsdata'  
  user='receiver'  
  aclgrp='PROD'  
  prompt=yes;
```

Note: Password values are case sensitive. That is, if the SAS Scalable Performance Data (SPD) Server administrator assigns a lowercase value, you must enter the password value in lowercase.

CHNGPASS=

Specifies to prompt a SAS Scalable Performance Data (SPD) Server user for a change of password. If ACL file security is enabled, the SAS Scalable Performance Data (SPD) Server validates the old/new password against its user ID table.

Syntax

```
CHNGPASS= YES | NO
```

Use the following arguments:

YES

prompts for a change of the SAS Scalable Performance Data (SPD) Server user password.

NO

suppresses a prompt for a change of the SAS Scalable Performance Data (SPD) Server user password. This is the default.

Example

Specify a prompt to change the password of SAS Scalable Performance Data (SPD) Server user TEMPHIRE.

```
LIBNAME mylib sasspds 'spdsdata'  
  user='temphire'  
  password='whizbang'  
  chngpass=yes;
```

Note: If you are using LDAP user authentication, and create a user connection that uses the CHNGPASS= LIBNAME option, the user password will not be changed. If you are using LDAP authentication and want to change a user password, follow the operating system procedures to change a user password, and check with your LDAP server administrator to ensure that the LDAP database also records password changes.

For more information on LDAP user authentication, see the reference help chapter in the SAS Scalable Performance Data (SPD) Server User's Guide, "Managing SAS Scalable Performance Data (SPD) Server Passwords, Users, and Table ACLs."

NEWPASSWORD= or NEWPASSWD=

Specifies a new password for a SAS Scalable Performance Data (SPD) Server client user. If ACL file security is enabled, the SAS Scalable Performance Data (SPD) Server validates the old/new password against its user ID table.

Syntax

```
NEWPASSWORD= 'newpassword'  
NEWPASSWD= 'newpassword'
```

Use the following argument:

'newpassword'

is the new password of a SAS Scalable Performance Data (SPD) Server client user. The password, visible in a SAS program, is encrypted in the SAS log file.

Example

Specify a new password **rambo** for SAS Scalable Performance Data (SPD) Server client user RECEIVER.

```
LIBNAME mylib sasspds 'spdsdata'  
  user='receiver'  
  password='whizbang'  
  newpassword='rambo';
```

Note: If you are using LDAP user authentication, and create a user connection that uses the [NEWPASSWORD=](#) LIBNAME option, the user password will not be changed. If you are using LDAP authentication and want to change a user password, follow the operating system procedures to change a user password, and check with your

LDAP server administrator to ensure that the LDAP database also records password changes.

For more information on LDAP user authentication, see the reference help chapter in the SAS Scalable Performance Data (SPD) Server User's Guide, "Managing SAS Scalable Performance Data (SPD) Server Passwords, Users, and Table ACLs."

PASSWORD= or PASSWD=

Specifies the SAS Scalable Performance Data (SPD) Server password of a SAS Scalable Performance Data (SPD) Server client user. If ACL file security is enabled, the SAS Scalable Performance Data (SPD) Server validates the password against its user ID table.

Syntax

```
PASSWORD='password'  
PASSWD='password'
```

Use the following argument:

'password'

is the case-sensitive password of a SAS Scalable Performance Data (SPD) Server client user. The password, visible in a SAS program, is encrypted in the SAS log file.

Example

Specify the password **whizbang** for SAS Scalable Performance Data (SPD) Server client user SPDSUSER.

```
LIBNAME mylib sasspds 'spdsdata'  
server=kaboom.5200  
user='spdsuser'  
password='whizbang';
```

Options

SAS Scalable Performance Data (SPD) Server 4.4 supports the integration of the SAS 9.1.3 PROC PWENCODE. This permits scripts to be generated that do not explicitly contain secure passwords that could easily be used without authorization. You must run PROC PWENCODE in Base SAS to enable the usage of script password encoding within SAS Scalable Performance Data (SPD) Server 4.4. See the Base SAS documentation for detailed instruction on running PROC PWENCODE for use with SAS Scalable Performance Data (SPD) Server 4.4.

The example below shows a SAS Scalable Performance Data (SPD) Server LIBNAME statement that utilizes the password encoding option:

```
LIBNAME mylib sasspds 'spdsdata'  
server=kaboom.5200  
user='spdsuser'  
password='{sas001}c3BkczeYmW==';
```

PROMPT=

Specifies to prompt a SAS Scalable Performance Data (SPD) Server user for a password. If ACL file security is enabled, the SAS Scalable Performance Data (SPD) Server validates the password against its user ID table.

Syntax

```
PROMPT= YES | NO
```

Use the following arguments:

YES

prompts a SAS Scalable Performance Data (SPD) Server user for a password.

NO

suppresses a prompt for a password.

Example

Specify a prompt for the password of SAS Scalable Performance Data (SPD) Server user BIGWHIG.

```
LIBNAME mylib sasspds 'spdsdata'  
  user='bigwhig'  
  prompt=yes;
```

USER=

Specifies the ID of a SAS Scalable Performance Data (SPD) Server client user. If ACL file security is enabled, the SAS Scalable Performance Data (SPD) Server validates the ID against its user ID table. (The SAS Scalable Performance Data (SPD) Server user ID defaults to the SAS process user ID if it is available; that is, when the client is not a Windows client.)

Syntax

```
USER= 'username'
```

Use the following argument:

'username'

is the ID of a SAS Scalable Performance Data (SPD) Server client user.

Example

Specify the identifier SPDSUSER for a SAS Scalable Performance Data (SPD) Server client user.

```
LIBNAME mylib sasspds 'spdsdata'  
  user='spdsuser'  
  prompt=yes;
```

Options to Specify Implicit SQL Pass-Through

IP=YES

This is an abbreviated specification which replaces the more verbose PASSTHRU= option. The IP=YES option draws on information specified in the LIBNAME declaration. The IP=YES option specifies an implicit SQL pass-through connection for a single user to a specified domain and server during a given SAS Scalable Performance Data (SPD) Server session.

Syntax

```
LIBNAME BOAF sasspds 'BOAF'  
  server=kaboom.5200  
  user='rcnye'  
  password='*****'  
  IP=YES ;
```

PASSTHRU=

This older and more verbose specification for IP=YES is still supported. It specifies an implicit SQL pass-through connection for a single user to a specified domain and server during a given SAS Scalable Performance Data (SPD) Server session.

Syntax

```
PASSTHRU=<'dbq=<"SAS-data-library">  
<SPD Server-options>  
user=<'UserID'>  
password=<'password'> ;
```

Use the following arguments:

DBQ=libname-domain (required)

Specifies the primary SAS Scalable Performance Data (SPD) Server LIBNAME domain for the SQL pass-through connection.

The name that you specify is identical to the LIBNAME domain name that you used when making a SAS LIBNAME assignment to **sasspds**. Use single or double quotes around the specified value.

SPD Server-options

one or more SAS Scalable Performance Data (SPD) Server options.

USER=SAS Scalable Performance Data (SPD) Server user ID (required on Windows but not UNIX)

Specifies a SAS Scalable Performance Data (SPD) Server user ID to access a SAS Scalable Performance Data (SPD) Server SQL Server. Use single or double quotes around the specified value.

PASSWORD=password (required, or use PROMPT=YES, unless USER='anonymou')

Specifies a SAS Scalable Performance Data (SPD) Server user ID password to access a SAS Scalable Performance Data (SPD) Server. (This value is case sensitive.)

Example:

The following is a LIBNAME statement that specifies the implicit SQL pass-through option for user rcnye, using a LIBREF to connect to the domain named 'BOAF' on the server named 'Kaboom' on port 5200:

```
LIBNAME BOAF sasspds 'BOAF'
  server=kaboom.5200
  user='rcnye'
  password='*****'

PASSTHRU= '
dbq="BOAF"
server=kaboom.5200
user="rcnye"
password="*****" ' ;
```

Options

SAS Scalable Performance Data (SPD) Server 4.4 supports the integration of the SAS 9.1.3 PROC PWENCODE. This permits scripts to be generated that do not explicitly contain secure passwords that could easily be used without authorization. You must run PROC PWENCODE in Base SAS to enable the usage of script password encoding within SAS Scalable Performance Data (SPD) Server 4.4. See the Base SAS documentation for detailed instruction on running PROC PWENCODE for use with SAS Scalable Performance Data (SPD) Server 4.4.

The example below shows a SAS Scalable Performance Data (SPD) Server LIBNAME statement that utilizes the password encoding option:

```
LIBNAME mylib sasspds 'spdsdata'
  server=kaboom.5200
  user='spdsuser'
  password='{sas001}c3BkczyEYMw=='

PASSTHRU= '
dbq="spdsdata"
server=kaboom.5200
user="spdsuser"
password="{sas001}c3BkczyEYMw==" ' ;
```

Options to Specify File Paths for Table Storage

SAS Scalable Performance Data (SPD) Server strongly recommends that your site administrator defines SAS Scalable Performance Data (SPD) Server domain options in the SAS Scalable Performance Data (SPD) Server `libnames.parm` configuration file. However, in unusual cases, such as the SAS Scalable Performance Data (SPD) Server administrator being temporarily unavailable, the following four LIBNAME options can be issued by a SAS Scalable Performance Data (SPD) Server user to define domains and table file storage paths.

CREATE=

Creates the primary directory for a SAS Scalable Performance Data (SPD) Server domain if it does not already

exist.

Syntax

CREATE=YES | NO

Use the following arguments:

YES

creates the primary directory if it does not already exist.

NO

fails the LIBNAME assignment if the primary directory does not already exist. This is the default.

Description

a SAS Scalable Performance Data (SPD) Server administrator defines the primary directory for the SAS Scalable Performance Data (SPD) Server domain in the LIBNAME parameter file. If CREATE= is YES, the software creates the directory (primary file system) in the event that a SAS Scalable Performance Data (SPD) Server administrator forgets to create it.

DATAPATH=

The SAS Scalable Performance Data (SPD) Server administrator for your site should use options in the SAS Scalable Performance Data (SPD) Server **libnames.parm** configuration file to define SAS Scalable Performance Data (SPD) Server domain options. However, if the SAS Scalable Performance Data (SPD) Server is temporarily unavailable, the following LIBNAME option can be issued by a SAS Scalable Performance Data (SPD) Server user to specify a list of initial or overflow paths to store data (.dpf) file partitions for a SAS Scalable Performance Data (SPD) Server table.

Syntax

DATAPATH=('filesystem' 'filesystem'...)

Use the following argument:

'filesystem'

is a directory path for UNIX or Windows.

Example

Create partitions as needed by cycling through the directories specified, DATAFLOW1 directory on DISK1 and DATAFLOW2 directory on DISK2.

```
LIBNAME mylib sasspds 'spdsdata'  
      datapath=('/disk1/dataflow1'  
              '/disk2/dataflow2');
```

INDEXPATH=

The SAS Scalable Performance Data (SPD) Server administrator for your site should use options in the SAS Scalable Performance Data (SPD) Server `libnames.parm` configuration file to define SAS Scalable Performance Data (SPD) Server domain options. However, if the SAS Scalable Performance Data (SPD) Server is temporarily unavailable, the following LIBNAME option can be issued by a SAS Scalable Performance Data (SPD) Server user to specify a list of initial or overflow paths to store index (. hbx), (. idx), and (. aux) file partitions associated with a SAS Scalable Performance Data (SPD) Server table.

Syntax

```
INDEXPATH=('filesystem' 'filesystem'...)
```

Use the following argument:

filesystem'

is a directory path for UNIX or Windows.

Example

Create index file partitions as needed using the directories specified, IDXFLOW1 directory on DISK1 and IDXFLOW2 directory on DISK2.

```
LIBNAME mylib sasspds 'spdsdata'  
indexpath=( '/disk1/idxflow1'  
            '/disk2/idxflow2');
```

METAPATH=

The SAS Scalable Performance Data (SPD) Server administrator for your site should use options in the SAS Scalable Performance Data (SPD) Server `libnames.parm` configuration file to define SAS Scalable Performance Data (SPD) Server domain options. However, if the SAS Scalable Performance Data (SPD) Server is temporarily unavailable, the following LIBNAME option can be issued by a SAS Scalable Performance Data (SPD) Server user to specify a list of overflow paths to store metadata (. mdf) file partitions for a SAS Scalable Performance Data (SPD) Server table.

Syntax

```
METAPATH=( 'filesystem' 'filesystem'...)
```

Use the following argument:

filesystem'

is a directory path for UNIX or Windows.

Example

Create overflow metadata file partitions as needed using the directories specified, METAFLOW1 directory on DISK1 and METAFLOW2 directory on DISK2.

```
LIBNAME mylib sasspds 'spdsdata'  
  metapath=( '/disk1/metaflow1'  
             '/disk2/metaflow2');
```

Options for Access Control Lists (ACLs)

ACLSPECIAL=

Grants special access to SAS Scalable Performance Data (SPD) Server resources in the LIBNAME domain to a SAS Scalable Performance Data (SPD) Server user. The SAS Scalable Performance Data (SPD) Server user must also be defined as 'special' by the SAS Scalable Performance Data (SPD) Server administrator.

Syntax

```
ACLSPECIAL=YES | NO
```

Use the following arguments:

YES

grants special access (read, write, alter, and control permission) to all SAS Scalable Performance Data (SPD) Server resources in the domain.

NO

denies special access (read, write, alter, and control permission) to all SAS Scalable Performance Data (SPD) Server resources in the domain.

Description

Grants special privileges to all SAS Scalable Performance Data (SPD) Server tables and associated indexes in the LIBNAME domain. The special privileges, (read, write, alter, and control permissions), override normal ACL restrictions only if the SAS Scalable Performance Data (SPD) Server administrator defines the user as 'special' in the user ID table.

Example

Grant special privileges to THEBOSS allowing him to read, write, alter, and control all tables in the CONVERSION_AREA domain. (The SAS Scalable Performance Data (SPD) Server administrator has defined THEBOSS as 'special'.)

```
LIBNAME mydatalib sasspds 'conversion_area'  
  server=husky.5105  
  user='theboss'  
  prompt=yes  
  aclspecial=yes ;
```

Options for a Client and Server Running on the Same UNIX Machine

NETCOMP=

Compresses the data stream for a SAS Scalable Performance Data (SPD) Server network packet.

Syntax

```
NETCOMP=YES | NO
```

Use the following arguments:

YES

sends compressed data in a SAS Scalable Performance Data (SPD) Server network packet.

NO

sends uncompressed data in a SAS Scalable Performance Data (SPD) Server network packet.

Description

Normally, data compression for inter-process transfers is recommended. However, for a client and server process on the same machine -- with UNIXDOMAIN=YES -- turning off compression can improve performance. You should examine NETCOMP together with UNIXDOMAIN and NETPACKSIZE for both client and server on the same machine.

Example

Specify to turn off compression of the data stream.

```
LIBNAME mylib sasspds 'test_area'  
netcomp=no;
```

UNIXDOMAIN=

Specifies the use of UNIX domain sockets for data communication between a SAS Scalable Performance Data (SPD) Server and client process **running on the same machine**. (Not available in Windows.)

Syntax

```
UNIXDOMAIN=YES | NO
```

Use the following arguments:

YES

uses AF_UNIX domain sockets for client/server data communication.

NO

uses the default AF_INET domain sockets for client/server data communication.

Description

When UNIXDOMAIN=YES, SAS Scalable Performance Data (SPD) Server uses AF_UNIX domain sockets rather than the customary AF_INET domain sockets for data communication. AF_UNIX sockets typically are much faster and greatly enhance performance but are only possible for cases where client and server are running on the same machine. You should also examine NETCOMP and NETPACKSIZE parameters for possible use to enhance performance in conjunction with UNIXDOMAIN.

Example

You find that using the AF_UNIX sockets for your session that is running on the same machine as the SAS Scalable Performance Data (SPD) Server is not faster. Configure SAS Scalable Performance Data (SPD) Server to use the default AF_INET sockets instead.

```
LIBNAME mylib sasspds 'test_area'  
    unixdomain=no;
```

Note: If you are running SAS Scalable Performance Data (SPD) Server 4.4 or later, and the client and server are both running UNIX, SAS Scalable Performance Data (SPD) Server automatically detects UNIX domain sockets. In such cases, it is not necessary to specify the UNIXDOMAIN parameter for optimum performance.

Options for Other Functions

BYSORT=

Specifies an implicit automatic SAS Scalable Performance Data (SPD) Server sort for a BY clause.

Syntax

```
BYSORT=YES | NO
```

Use the following arguments:

YES

performs an implicit sort for a BY clause. This is the default.

NO

does not perform an implicit sort for a BY clause.

Description

Where Base SAS software requires an explicit sort statement (PROC SORT) to sort SAS data, by default, SAS Scalable Performance Data (SPD) Server performs a sort whenever it encounters a BY clause. If the value of the BYSORT= option is NO, the SAS Scalable Performance Data (SPD) Server software performs the same as the base SAS engine.

Example 1

Specify to turn off implicit SAS Scalable Performance Data (SPD) Server sorts for the session.

```

LIBNAME mydatalib sasspds 'conversion_area'
  server=husky.5105
  user='siteusr1'
  prompt=yes
  bysort=no ;

data mydatalib.old_autos;
  input
    year $4.
    @6 manufacturer $12.
    model $10.
    body_style $5.
    engine_liters
    @39 transmission_type $1.
    @41 exterior_color $10.
    options $10.
    mileage condition ;
  datalines ;

1971 Buick      Skylark   conv   5.8   A   yellow   00000001 143000 2
1982 Ford      Fiesta    hatch  1.2   M   silver   00000001  70000 3
1975 Lancia    Beta      2door  1.8   M   dk blue  00000010  80000 4
1966 Oldsmobile Toronado  2door  7.0   A   black    11000010 110000 3
1969 Ford      Mustang   sptrf  7.1   M   red      00000111 125000 3
;

PROC PRINT data=mydatalib.old_autos;
  by model;
run;

```

In this program, the PRINT procedure will return an error message because the table MYDATALIB.OLD_AUTOS is not sorted.

Example 2

Turn off implicit SAS Scalable Performance Data (SPD) Server sorts with the LIBNAME option, but specify a server sort for the table MYDATALIB.OLD_AUTOS using the BYSORT table option.

```

PROC PRINT data=mydatalib.old_autos
  (bysort=yes);
  by model;
run;

```

DISCONNECT=

Specifies when to close the network connections between the SAS client and the SAS Scalable Performance Data (SPD) Server. This can be done either when all SAS Scalable Performance Data (SPD) Server librefs are cleared or when the SAS client session ends.

Syntax

DISCONNECT=YES | NO

Use the following arguments:

YES

closes network connections between the SAS client and SAS Scalable Performance Data (SPD) Server when all SAS Scalable Performance Data (SPD) Server librefs are cleared.

NO

closes network connections between the SAS client and SAS Scalable Performance Data (SPD) Server only when the SAS session ends. This is the default.

Description

When the DISCONNECT= option is NO, the network connections between the SAS client and the SAS Scalable Performance Data (SPD) Server are closed when the current SAS session ends. When network connections are active, the user can issue successive librefs to the same SAS Scalable Performance Data (SPD) Server more efficiently.

When the DISCONNECT= option is YES, the server connection closes after all SAS Scalable Performance Data (SPD) Server librefs are cleared. Assuming a user does not issue a subsequent LIBNAME statement, closing the connection frees resources. For example, a SAS job or program accesses a SAS Scalable Performance Data (SPD) Server table at the beginning of a job but performs remaining processing locally. In this situation, closing the network connection after clearing all librefs frees both SAS and SAS Scalable Performance Data (SPD) Server file descriptors, machine memory, and TCP/IP resources.

Note: Unless this option is used with the *initial* LIBNAME engine statement for the SAS Scalable Performance Data (SPD) Server session, it has no effect.

Example

Specify for SAS Scalable Performance Data (SPD) Server to close the network connections after you clear your librefs, rather than at the end of your SAS session.

```
LIBNAME spud sasspds 'potatoes'
    disconnect=yes
    server=husky.5105
    user='siteusr1'
    prompt=yes ;

data local;
    set spud.idaho ;
run ;

/* Clear the libref SPUD so SPD Server will close the server */
/* connection - Do the rest of the SAS processing locally */

LIBNAME spud clear;

/* The rest of the program follows */
```


ENDOBS=

Specifies the end row (observation) number in a user-defined range for processing.

Syntax

```
ENDOBS=n
```

Use the following argument:

n
is the number of the end row.

Description

By default, SAS Scalable Performance Data (SPD) Server processes the entire table unless the user specifies a range of rows with the STARTOBS= and ENDOBS= options. If the STARTOBS= option is used without the ENDOBS= option, the implied value of ENDOBS= is the end of the table. When both options are used together, the value of ENDOBS= must be greater than STARTOBS=.

In contrast to the Base SAS software options FIRSTOBS= and OBS=, the STARTOBS= and ENDOBS= SAS Scalable Performance Data (SPD) Server options can be used for WHERE-clause processing in addition to table input operations.

Example 1

Specify for SAS Scalable Performance Data (SPD) Server to process only row numbers (observations) 200 — 500 while the LIBNAME is active.

```
LIBNAME mydatalib sasspds 'conversion_area'  
server=husky.5105  
user='siteusr1'  
prompt=yes  
startobs=200  
endobs=500;
```

LIBGEN=

The LIBGEN=YES option is used in explicit SQL connection statements. When you set LIBGEN= yes, SAS Scalable Performance Data (SPD) Server is configured to generate additional domain connections that enable you to perform SQL joins across different SAS Scalable Performance Data (SPD) Server domains.

Syntax

```
LIBGEN=YES
```

Description

You should specify the LIBGEN=YES option in the explicit SQL LIBNAME connection. You cannot specify the LIBGEN=YES option setting without first creating a LIBNAME connections to the domain.

Examples

The two code examples that follow both perform the same task. Both examples use explicit SQL to join two tables from different domains. The first example uses execute connection statements to facilitate joining the tables from separate domains. The second example uses the LIBGEN=YES option to perform the same join without having to issue the extra execute connection statements.

SQL Without LIBGEN=YES

```
/* The example code without LIBGEN=YES */
/* must issue execute connection      */
/* statements to access tables that   */
/* reside in two different domains.   */

LIBNAME path1 sasspds 'path1'
  server=boxer.5140
  ip=YES
  user='anonymous' ;

LIBNAME path2 sasspds 'path2'
  server=boxer.5140
  ip=YES
  user='anonymous' ;

DATA path1.table1
  (keep=i table1)
path2.table2
  (keep=i table2) ;

table1 = 'table1' ;
table2 = 'table2' ;

do i = 1 to 10 ;
  output ;
end ;
run ;

PROC SQL ;
CONNECT to sasspds (
  dbq='Path1'
  server=boxer.5140
  user='anonymous') ;

/* Without LIBGEN=YES, you must make */
/* two execute connect statements.   */

execute (LIBREF path1 engopt="dbq='path1'")
  by sasspds;
execute (LIBREF path2 engopt="dbq='path2'")
  by sasspds;

execute
  (create table table4 as
```

```

select *
from
  path1.table1 a,
  path2.table2 b
where a.i = b.i)
by sasspds ;

disconnect from sasspds ;

quit ;

```

SQL With LIBGEN=YES

```

/* The example code that uses LIBGEN=YES */
/* can join the tables from two different */
/* domains in a more simple manner. */

LIBNAME path1 sasspds 'path1'
  server=boxer.5140
  LIBGEN=YES
  ip=YES
  user='anonymous' ;

LIBNAME path2 sasspds 'path2'
  server=boxer.5140
  LIBGEN=YES
  ip=YES
  user='anonymous' ;

DATA path1.table1
  (keep=i table1)
path2.table2
  (keep=i table2) ;

table1 = 'table1' ;
table2 = 'table2' ;

do i = 1 to 10 ;
  output ;
end ;
run ;

PROC SQL ;
CONNECT to sasspds (
  dbq='Path1'
  server=boxer.5140
  user='anonymous') ;

/* Syntax used with LIBGEN=YES option */

execute
  (create table table4 as

```

```
select *
from
  path1.table1 a,
  path2.table2 b
where a.i = b.i)
by sasspds ;

disconnect from sasspds ;

quit ;
```

LOCKING=

Overview of Record-Level Locking

Record-level locking is a SAS Scalable Performance Data (SPD) Server feature that allows multiple users concurrent read and write access to SAS Scalable Performance Data (SPD) Server tables while maintaining the integrity of the table contents. When record-level locking is enabled, users can insert, append, delete, and update the contents of a SAS Scalable Performance Data (SPD) Server table while performing concurrent reads on the table. When a client enables record-level locking, the client connects to the single SAS Scalable Performance Data (SPD) Server record-level locking proxy process. When record-level locking is not enabled, clients connect to separate SAS Scalable Performance Data (SPD) Server user proxy processes for each LIBNAME connection to a domain.

Record-Level Locking Details

Record-level locking is enabled when a SAS Scalable Performance Data (SPD) Server client specifies the LOCKING=YES LIBNAME option to the client's LIBNAME connection statement. All subsequent operations on the given LIBNAME domain will employ record-level locking. The primary use of record-level locking is to allow multiple clients or parallel operations from the same client to have both read and write access to the same SAS Scalable Performance Data (SPD) Server table resource. If record-level locking is not enabled, then any write operation (update, append, insert, or delete) to a SAS Scalable Performance Data (SPD) Server table requires exclusive access to the resource, or else a member lock failure error occurs. Operations that affect metadata, such as creating or deleting indexes, renaming variables, and renaming tables require exclusive access to the resource, whether record-level locking is enabled or not. These types of operations will report a member lock failure error when with record-level locking is enabled, but exclusive access is not available.

Record-level locking must be enabled in SAS Scalable Performance Data (SPD) Server before a SAS client can use the CNTLEV=REC table option in their SAS program to access SAS Scalable Performance Data (SPD) Server tables. Record-level locking enforces SAS style record-level integrity across multiple clients, so clients are guaranteed that an observation will not change during a multi-phased read or write operation on the specified observation. Record-level locking will allow multiple concurrent update access to a single SAS Scalable Performance Data (SPD) Server table, but it will deny concurrent access to the specified observation within the table.

When a SAS Scalable Performance Data (SPD) Server client establishes a LIBNAME connection to a domain with record-level locking enabled, it connects using the single record-level locking proxy process. There is only one record-level locking proxy process per SAS Scalable Performance Data (SPD) Server. All SAS Scalable Performance Data (SPD) Server clients that use record-level locking connections are processed through the record-level locking proxy process. If there are a large number of record-level locking connections, there may be some contention for process resources between the clients. The record-level locking proxy process is

a single point of failure for all these connections, so care should be taken when you use record-level locking to update critical data.

When you append or insert new rows into a table with defined indexes, the table updates are processed more sequentially through the record-level locking proxy process than they would be through the SPD user proxy processes. The performance of record-level locking will probably be less than the performance that can be obtained without record-level locking enabled for these types of operations. The standard member-level locking that is used in SPD user proxy processes allows for more parallel processing when doing table append or insert operations.

Record-level locking is not supported for operations on tables that use dynamic clusters.

Syntax

```
LOCKING=YES | NO
```

Default: NO

Use the following arguments:

YES

enables record sharing mode.

NO

disables record sharing mode.

Example

```
LIBNAME testrl sasspds 'tmp'  
server=serverNode.port  
user='anonymous'  
locking=YES ;
```

STARTOBS=

Specifies the start row (observation) number in a user-defined range for processing.

Syntax

```
STARTOBS=n
```

Use the following arguments:

n

is the number of the start row.

Description

By default, SAS Scalable Performance Data (SPD) Server processes the entire table unless the user specifies a range of rows with the options, STARTOBS= and ENDOBS=. If the ENDOBS= option is used without the STARTOBS= option, the implied value of STARTOBS= is 1. When both options are used together, the value of

STARTOBS= must be less than the value of ENDOBS.=

In contrast to the Base SAS software options FIRSTOBS= and OBS=, the STARTOBS= and ENDOBS= SAS Scalable Performance Data (SPD) Server options can be used for WHERE-clause processing in addition to table input operations.

Example

Specify for SAS Scalable Performance Data (SPD) Server to process only row numbers (observations) 200 — 500 while the LIBNAME is active.

```
LIBNAME mydatalib sasspds 'conversion_area'  
  server=husky.5105  
  user='siteusr1'  
  prompt=yes  
  startobs=200  
  endobs=500;
```

TEMP=

Controls the creation of a temporary LIBNAME domain for this LIBNAME assignment.

Syntax

```
TEMP=YES | NO
```

Default: NO

Use the following arguments:

YES

creates a temporary LIBNAME domain for the LIBNAME assignment.

NO

does not create a temporary LIBNAME domain.

Description

Use this option to create temporary LIBNAME domains that exist for the duration of the LIBNAME assignment. The TEMP (temporary) domains are analogous to SAS WORK libraries.

To create a temporary LIBNAME domain, use TEMP=YES. Any data objects, tables, catalogs, or utility files that are created in the TEMP=YES temporary domain are automatically deleted when you end the SAS session. This functions similarly to a SAS WORK library. (Note: The temporary domain is created as a subdirectory of the directory specified as the library domain.)

Example 1

Create a LIBNAME domain to use for temporary storage during your SAS session.

```
LIBNAME mydatalib sasspds 'conversion_area'  
  server=kaboom.5191  
  user='siteusr1'  
  prompt=yes  
  temp=yes ;
```

TRUNCWARN=

Suppresses hard failure on NLS transcoding overflow and character mapping errors.

Syntax

```
TRUNCWARN=YES | NO
```

Default: NO

Description

When using the TRUNCWARN=YES LIBNAME option, data integrity may be compromised because significant characters can be lost in this configuration. The default setting is NO, which causes hard read/write stops when transcode overflow or mapping errors are encountered. When TRUNCWARN=YES, and an overflow or character mapping error occurs, a warning is posted to the SAS log at data set close time if overflow occurs, but the data overflow is lost.

WORKPATH=

I/O contention can occur when many SAS Scalable Performance Data (SPD) Server users or SAS Scalable Performance Data (SPD) Server jobs perform heavy processing that uses the same workpath. The WORKPATH= option permits users to specify an alternate workpath that utility files (such as index builds and sorting files) can use. Specifying an alternate workpath can relieve I/O contention issues when other users are running heavy processing jobs at the same time.

A properly configured workpath directs I/O from utility operations to a separate disk. Mapping the utility file work to a separate disk using the WORKPATH= option avoids conflict with other jobs that use a default workpath that is specified in the spdsserv.parm configuration file.

Using the optional WORKPATH= specification to direct utility file operations to a separate disk increases the overall I/O through-put for the utility files and speeds up the server performance as well.

Syntax

```
WORKPATH=('path-specification');
```

Example

Two SAS Scalable Performance Data (SPD) Server power users perform heavy index creation and are creating heavy I/O contention on the default workpath that is defined in the spdsserv.parm configuration file:

```
workpath=('workspace1')
```

Both users override the default workpath by using the alternate WORKPATH= specification when issuing the LIBNAME statements in their jobs:

User 1 LIBNAME statement:

```
LIBNAME domain-name sasspds "domain-name"  
      server=host-name.port-number  
      user='user1'  
  
workpath=('/bigdisk/spdsmgr/workpath1') ;
```

User 2 LIBNAME statement:

```
LIBNAME domain-name sasspds "domain-name"  
      server=host-name.port-number  
      user='user2'  
  
workpath=('/bigdisk/spdsmgr/workpath2') ;
```

All SAS Scalable Performance Data (SPD) Server jobs by other users continue to use the default workpath specification that is declared in spdsserv.parm

The LIBNAMEs.parm configuration file also accepts alternate WORKPATH= specifications for each domain.

Copyright © 2008 by SAS Institute Inc., Cary, NC, USA. All rights reserved.

SAS Scalable Performance Data (SPD) Server Table Options

Contents

- [Introduction](#)
- [Option for Compatibility with Base SAS Software](#)
 - [SYNCADD=](#)
- [Options That Affect Disk Space](#)
 - [ASYNCINDEX=](#)
 - [COMPRESS=](#)
 - [PARTSIZE=](#)
- [Options to Enhance Performance](#)
 - [BYNOEQUALS=](#)
 - [IOBLOCKSIZE=](#)
 - [NETPACKSIZE=](#)
 - [SEGFSIZE=](#)
- [Option to Test Performance](#)
 - [NOINDEX=](#)
- [Options for WHERE Clause Evaluations](#)
 - [MINMAXVARLIST=](#)
 - [THREADNUM=](#)
 - [WHEREINDEX=](#)
- [Options for Other Functions](#)
 - [BYSORT=](#)
 - [ENDOBS=](#)
 - [STARTOBS=](#)
 - [UNIQUESAVE=](#)
 - [VERBOSE=](#)
- [Options for Security](#)
 - [ENCRYPT=](#)

[Introduction](#)

All SAS users who use LIBNAME access to SAS Scalable Performance Data (SPD) Server should read this chapter. Most table options also work in SQL pass-through statements.

This chapter presents reference information for the SAS Scalable Performance Data (SPD) Server table options. To specify a table option with LIBNAME access, place the option value in parentheses after the table name. The option value then specifies processing that applies only to that table. To specify a table option with pass-through access, place the option value in brackets after the table name. The option value then specifies processing that applies only to that table. The SAS Scalable Performance Data (SPD) Server table options that follow are grouped by the function of their default value.

When using the options in this chapter, remember that if a table option is used subsequent to a LIBNAME option or macro variable, the value of the table option takes precedence.

[Option for Compatibility with Base SAS Software](#)

[SYNCADD=](#)

Specifies when appending to a table whether to apply a single or multiple rows at a time.

Syntax

```
SYNCADD=YES|NO
```

Default

```
NO
```

Corresponding Macro Variable

```
SPDSSADD
```

Related Table Option

```
UNIQUESAVE=
```

Use the following arguments:

YES

imitates the behavior of the base SAS engine, applying a single row at a time (synchronously).

NO

appends multiple rows at a time (asynchronously).

Description

When SYNCADD= is YES, processing performance is slower. Use this setting only to force the server's append processing to be compatible with Base SAS software processing. That is, when the server encounters a row with a non-unique value, to abort the append operation, back out the transactions just added, and leave the original table on disk.

Example

In this example, when executing the first INSERT statement, PROC SQL permits insertion of the values 'rollback1' and 'rollback2' because the row additions to table A are performed asynchronously. PROC SQL does not get the true completion status at the time it adds a row.

When executing the second INSERT statement, PROC SQL performs a rollback on the INSERT, upon encountering the Add error on 'nonunique', and deletes the rows 'rollback3' and 'rollback4'.

```
data a;
  input z $ 1-20 x y;
  list;

  datalines;

one          1 10
two          2 20
three        3 30
four         4 40
five         5 50
;

PROC SQL sortseq=ascii exec noerrorstop;
create unique index comp on a (x, y);
```

```
insert into a
  values('rollback1', -80, -80)
  values('rollback2',-90, -90)
  values('nonunique', 2, 20);

insert into a(syncadd=yes)
  set z='rollback3', x=-60, y=-60
  set z='rollback4', x=-70, y=-70
  set z='nonunique', x=2, y=20;
quit;
```

Options That Affect Disk Space

ASYNINDEX=

Specifies when creating multiple indexes on a SAS Scalable Performance Data (SPD) Server table whether to create the indexes in parallel.

Syntax

```
ASYNINDEX=YES|NO
```

Default

```
NO
```

Corresponding Macro Variable

```
SPDSIASY
```

Use the following arguments:

YES

creates the indexes in parallel.

NO

creates a single index at a time.

Description

The SAS Scalable Performance Data (SPD) Server can create multiple indexes for a table at the same time. To do this, it launches a single thread for each index created, then processes the threads simultaneously. Although creating indexes in parallel is much faster, the default for this option is NO. The reason is because parallel creation requires additional sort work space which may not be available.

For a complete description of the benefits and tradeoffs of creating multiple indexes in parallel, see the SAS Scalable Performance Data (SPD) Server User's Guide Help section on [SPDSIASY=](#).

Example

Since the disk work space required for parallel index creation is available, specify for SAS Scalable Performance Data (SPD) Server to create, in parallel, the X, Y, and COMP indexes for table A.

```
PROC DATASETS lib=mydatalib;
```

```
modify a(asyncindex=yes);
index create x;
index create y;
index create comp=(x y);
quit;
```

COMPRESS=

Compresses SAS Scalable Performance Data (SPD) Server tables on disk.

Syntax

COMPRESS=YES|NO

Default

NO

Use in Conjunction with Table Option

IOBLOCKSIZE=

Corresponding Macro Variable

SPDSDCMP

Use the following arguments:

YES

performs the run-length compression algorithm SPDSRLLC.

NO

performs no table compression.

Description

When COMPRESS= is assigned YES, SAS Scalable Performance Data (SPD) Server compresses newly created tables by 'blocks' based on the algorithm specified. To control the amount of compression, use the table option [IOBLOCKSIZE=](#). This option specifies the number of rows that you want to store in the block.

Note: Once a compressed table is created, you cannot change its block size. To resize the block, you must PROC COPY the table to a new table, setting IOBLOCKSIZE= to the block size desired for the output table.

PARTSIZE=

Specifies the size of a SAS Scalable Performance Data (SPD) Server table partition.

Syntax

PARTSIZE=*n*

Default

16 Megabytes

Corresponding Macro Variable

SPDSSIZE=

Affected by LIBNAME option

DATAPATH=

Use the following argument:

n
is the size of the partition in megabytes.

Description

Specifying PARTSIZE= forces the software to partition (split) SAS Scalable Performance Data (SPD) Server tables at the given size. The actual size is computed to accommodate the largest number of rows that will fit in the specified size of *n* Mbytes.

Use this option to improve performance of WHERE Clause evaluation on non-indexed table columns and on SQL GROUP_BY processing. By splitting the data portion of a Scalable Platform Data Server table at fixed-sized intervals, the software can introduce a high degree of scalability for these operations. The reason: it can launch threads in parallel to perform the evaluation on different partitions of the table, without the threat of file access contention between the threads. There is, however, a price for the table splits: an increased number of files, which are required to store the rows of the table.

Ultimately, scalability limits using PARTSIZE= depend on how you structure [DATAPATH=](#), a LIBNAME option discussed in the documentation on [Scalable Performance Data Server LIBNAME Options](#). Specifically, the limits depend on how you configure and spread the DATAPATH= file systems across striped volumes. You should spread each individual volume's striping configuration across multiple disk controllers/SCSI channels in the disk storage array. The goal for the configuration is, at the hardware level, to maximize parallelism during data retrieval.

The PARTSIZE= specification is limited by MINPARTSIZE=, a SAS Scalable Performance Data (SPD) Server parameter maintained by the SAS Scalable Performance Data (SPD) Server administrator. MINPARTSIZE= ensures that an over-zealous SAS user does not create arbitrarily small partitions, thereby generating a large number of files. The default for MINPARTSIZE= is 16 Mbytes and probably should not be lowered much beyond this value.

Note: The PARTSIZE value for a table cannot be changed after a table is created. To change the PARTSIZE, you must PROC COPY the table and use a different PARTSIZE option setting on the new (output) table.

Example

Using PROC SQL, extract a set of rows from an existing table to create a non-indexed table with a partition size of 32 Mbytes in a SAS job:

```
PROC SQL;
create table SPDSCEN.HR80SPDS(partsize=32)
  as select
    state,
    age,
    sex,
```

```

hour89,
industry,
occup
from SPDSCEN.PRECS
where hour89 > 40;
quit;

```

Options to Enhance Performance

BYNOEQUALS=

Specifies the output order of table rows with identical values for the BY column.

Syntax

```
BYNOEQUALS=YES | NO
```

Use the following arguments:

YES

does not guarantee the output order of table rows with identical values in a BY clause.

NO

guarantees the output order of table rows with identical values in a BY clause will be the relative table position of the rows from the input table. This is the default.

Example

Specify for SAS Scalable Performance Data (SPD) Server in the ensuing BY-column operation to output rows with identical values in the key column randomly.

```

data sport.racquets(index=(string));
  input
    raqname $20.
    @22 weight
    @28 balance $2.
    @32 flex
    @36 gripsize
    @42 string $3.
    @47 price
    @55 instock;

  datalines;

Solo Junior          10.1  N   2  3.75  syn   50.00  6
Solo Lobber          11.3  N  10  5.5   syn  160.00  1
Solo Queensize      10.9  HH   6  5.0   syn  130.00  3
Solo Kingsize       13.1  HH   5  5.6   syn  140.00  3
;

data sport.racqbal(bynoequal=yes);
  set sport.racquets;
  by balance;
run;

```

IOBLOCKSIZE=

Specifies the number of rows in a block to be stored in or read from a SAS Scalable Performance Data (SPD) Server table.

Syntax

IOBLOCKSIZE=*n*

Default

4096

Use in Conjunction with

Macro Variable [SPDSDCMP=](#) or Table Options [COMPRESS=](#) or ENCRYPT= .

Use the following argument:

n
is the size of the block.

Description

The software reads and stores a server table in blocks. IOBLOCKSIZE= is useful on compressed or encrypted tables. SAS Scalable Performance Data (SPD) Server software does not use IOBLOCKSIZE= on noncompressed or nonencrypted tables.

For tables that you compress or encrypt, using either the option COMPRESS= or the macro variable SPDSDCMP=, the IOBLOCKSIZE= specification determines the number of rows to include in the block. The specification applies to block compression as well as data I/O to and from disk. The IOBLOCKSIZE= value affects the table's organization on disk.

When using SAS Scalable Performance Data (SPD) Server table compression or encryption, specify an IOBLOCKSIZE= value that complements how the data is to be accessed, sequentially or randomly. Sequential access or operations requiring full table scans favor a large block size, for example 64K. In contrast, random access favors a smaller block size, for example 8K.

Example

A huge company mailing list is processed sequentially. Specify a block size for compression that is optimal for sequential access.

```
/* IOblocksize set to 64K */  
  
data sport.maillist(ioblocksize=65536 compress=yes);  
  input  
    name $ 1-20  
    address $ 21-57  
    phoneno $ 58-69  
    sex $71;  
  
  datalines;  
  
Douglas, Mike          3256 Main St., Cary, NC 27511          919-444-5555 M  
Walters, Ann Marie    256 Evans Dr., Durham, NC 27707        919-324-6786 F  
Turner, Julia         709 Cedar Rd., Cary, NC 27513         919-555-9045 F  
Cashwell, Jack       567 Scott Ln., Chapel Hill, NC 27514  919-533-3845 M
```

NETPACKSIZE=

Specifies the size of the SAS Scalable Performance Data (SPD) Server network data packet.

Syntax

NETPACKSIZE=size-of-packet

Use the following argument:

size-of-packet

is the size of the network packet in bytes.

Description

This option controls the size of the buffer used for data transfer between SAS Scalable Performance Data (SPD) Server and a SAS client. The default is 32K bytes. The buffer size is relative to the size of a table row. It cannot be less than the size of a single row. Packet size must be equal to some multiple of the table rows. If it is not, SAS Scalable Performance Data (SPD) Server rounds up the size specified. For example, if the packet buffer size is 4096 bytes and the row size is 3072, the software rounds up the buffer size to 6144.

Select a packet size to complement the bandwidth of the network it must travel through. An optimum size will flow the data continuously without significant pauses between packets.

Example

Create a 12K buffer in the memory of the server to send three rows from MYTABLE in each network packet. (The row size in MYTABLE is 4K.)

```
data mylib.mytable (netpacksize=12288);
```

SEGSIZE=

Specifies the size of the segment for an index file associated with a SAS Scalable Performance Data (SPD) Server table.

Syntax

SEGSIZE=*number*

Use the following argument:

number

is the number of table rows to include in the index segment.

Description

The minimum SEGSIZE= value is 1024 table rows. The default value is 8192 table rows. The size of the index segment corresponds to the structure of the table and cannot be changed after the table is created.

Example

Specify a segment size of 64K for MYLIB.MYTABLE.

```
data mylib.mytable (segsz=65536);
```

Note: Tests show that increasing the size of the segment does not significantly increase performance.

Option to Test Performance

NOINDEX=

Specifies whether to use the table's indexes when processing WHERE Clauses.

Syntax

```
NOINDEX=YES|NO
```

Default

```
NO
```

Use the following arguments:

YES

ignores indexes when processing WHERE Clauses.

NO

uses indexes when processing WHERE Clauses.

Description

Set NOINDEX= to YES to test the effect of indexes on performance or for specific processing. *Do not use YES routinely for normal processing.*

Example

We created an index for the SEX column but decide to test whether it is necessary for our PROC PRINT processing. Specify for the server not to use the index.

```
data sport.maillist;
  input
    name $ 1-20
    address $ 21-57
    phoneno $ 58-69
    sex $71;

  datalines;
```

```
Douglas, Mike          3256 Main St., Cary, NC 27511      919-444-5555 M
Walters, Ann Marie    256 Evans Dr., Durham, NC 27707    919-324-6786 F
Turner, Julia         709 Cedar Rd., Cary, NC 27513     919-555-9045 F
Cashwell, Jack       567 Scott Ln., Chapel Hill, NC 27514 919-533-3845 M
Clark, John           9 Church St., Durham, NC 27705    919-324-0390 M
;
```

```

PROC DATASETS lib=sport nolist;
  modify maillist;
  index create sex;
quit;

/*Turn on the macro variable SPDSWDEB */
/* to show that the index is not used */
/* used during the table processing. */

%let spdswdeb=YES;

title "All Females from Current Mailing List";
PROC PRINT data=sport.maillist(noindex=yes);
where sex="F";
run;

```

Options for WHERE Clause Evaluations

MINMAXVARLIST=

Creates an index that documents the minimum and maximum values of specified variables. SAS Scalable Performance Data (SPD) Server WHERE Clause evaluations use MINMAXVARLIST= indexes to include or eliminate member tables in a SAS Scalable Performance Data (SPD) Server dynamic cluster table from SQL evaluation scans..

Syntax

MINMAXVARLIST=(*varname1*, *varname2*, ... , *varnameN*)

Use the following argument:

varname1, *varname2*, ... , *varname N*

are SAS Scalable Performance Data (SPD) Server table variable names.

Description

The primary purpose of the MIINMAXVARLIST= table option is for use with SAS Scalable Performance Data (SPD) Server [dynamic cluster tables](#) where specific members in the dynamic cluster contain a set or range of values, such as sales data for a given month. When a SAS Scalable Performance Data (SPD) Server SQL sub setting where- clause specifies specific months from a range of sales data, the WHERE planner checks the min/ max indexes. Based on the min/max index information, the SAS Scalable Performance Data (SPD) Server WHERE planner includes or eliminates member tables in the dynamic cluster for evaluation.

MINMAXVARLIST= uses the list of columns you submit to build an index. The MINMAXVARLIST= index contains only the minimum and maximum values for each column. The WHERE Clause planner uses the index to filter SQL predicates quickly, and to include or eliminate member tables belonging to the cluster table from the evaluation.

Although the MINMAXVARLIST= table option is primarily intended for use with dynamic clusters, it also works on standard SAS Scalable Performance Data (SPD) Server tables. MINMAXVARLIST= can help reduce the need to create many indexes on a table, which can save valuable resources and space.

Example

```

%let domain=path3 ;
%let host=kaboom ;
%let port=5201 ;

libname &domain sasspds "&domain"
      server=&host..&port
      user='anonymous' ;

/* Create three tables called */
/* xy1, xy2, and xy3.          */

data &domain..xy1(minmaxvarlist=(x y));
  do x = 1 to 10;
  do y = 1 to 3;
  output;
  end;
end;
run;

data &domain..xy2(minmaxvarlist=(x y));
  do x = 11 to 20;
  do y = 4 to 6 ;
  output;
  end;
end;
run;

data &domain..xy3(minmaxvarlist=(x y));
  do x = 21 to 30;
  do y = 7 to 9 ;
  output;
  end;
end;
run;

/* Create a dynamic cluster table */
/* called cluster_table out of    */
/* new tables xy1, xy2, and xy3   */

PROC SPDO library=&domain ;
  cluster create cluster_table
    mem=xy1
    mem=xy2
    mem=xy3
    maxslot=10;
quit;

/* Enable WHERE evaluation to see */
/* how the SQL planner selects     */
/* members from the cluster. Each */
/* member is evaluated using the   */
/* min-max index.                  */

%let SPDSWDEB=YES;

/* The first member has true rows */

PROC PRINT data=&domain..cluster_table ;

```

```

    where x eq 3
    and y eq 3;
run;

/* Examine the other tables */

PROC PRINT data=&domain..cluster_table ;
    where x eq 3
    and y eq 3 ;
run;

PROC PRINT data=&domain..cluster_table ;
    where x eq 3
    and y eq 3;
run;

PROC PRINT data=&domain..cluster_table ;
    where x between 1 and 10
    and y eq 3;
run;

PROC PRINT data=&domain..cluster_table ;
    where x between 11 and 30
    and y eq 8 ;
run;

/* Delete the dynamic cluster table. */

PROC DATASETS lib=&domain nolist;
    delete cluster_table ;
quit ;

```

THREADNUM=

Specifies the number of threads to be used for WHERE Clause evaluations.

Syntax

```
THREADNUM=n
```

Default

THREADNUM= is set equal to the value of the MAXWHTHREADS server parameter.

Used in Conjunction with SAS Scalable Performance Data (SPD) Server Parameter

```
MAXWHTHREADS
```

Corresponding Macro Variable

```
SPDSTCNT=
```

Use the following argument:

n

is the number of threads.

Description

THREADNUM= allows you to specify the thread count the SAS Scalable Performance Data (SPD) Server should use when performing a parallel WHERE Clause evaluation.

Use this option to explore scalability for WHERE Clause and GROUP_BY evaluations in non-production jobs. If you use this option for production jobs, you are likely to lower the level of parallelism that is applied to those clause evaluations.

THREADNUM= works in conjunction with MAXWHTHREADS, a configurable system parameter. MAXWHTHREADS imposes an upper limit on the consumption of system resources. The default value of MAXWHTHREADS is dependent on your operating system. Your SAS Scalable Performance Data (SPD) Server administrator can change the default value for MAXWHTHREADS.

If you do not use THREADNUM=, the software provides a default thread number, up to the value of MAXWHTHREADS as required. If you use THREADNUM=, the value that you specify is also constrained by the MAXWHTHREADS value.

The THREADNUM= value applies both to parallel table scans (EVAL2 strategy), parallel indexed evaluations (EVAL1 strategy), parallel BY-clause processing, and parallel GROUP_BY evaluations. The SAS Scalable Performance Data (SPD) Server User's Guide Help section on Optimizing SAS Scalable Performance Data (SPD) Server Performance, contains [more information on WHERE Clause evaluation](#).

Example

The SAS Scalable Performance Data (SPD) Server administrator set MAXWHTHREADS=128 in the SAS Scalable Performance Data (SPD) Server's parameter file. Explore the effects of parallelism on a given query by using the following SAS macro:

```
%macro dotest(maxthr);
%do nthr=1 %to &maxthr;
  data _null_;
    set SPDSZEN.PRECS(threadnum=&nthr);
    WHERE
      occup='022'
      and state in('37','03','06','36');
  run;
%mend dotest;
```

WHEREINDEX=

Specifies a list of indexes to exclude when making WHERE Clause evaluations.

Syntax

WHEREINDEX=(*name1 name2...*)

Use the following arguments:

(*name1 name2...*)

a list of index names that you wish to exclude from the WHERE planner.

Example

We have a table PRECS with indexes defined as follows:

```
PROC DATASETS lib=spdsцен;  
modify precs(bitindex=(hour89));  
index create  
  stser=(state serialno)  
  occind=(occup industry)  
  hour89;  
quit;
```

When evaluating the next query, we want the SAS Scalable Performance Data (SPD) Server to exclude from consideration indexes for both the STATE and HOUR89 columns.

In this case, we know that our AND combination of the predicates for the OCCUP and INDUSTRY columns will produce a very small yield. Few rows satisfy the respective predicates. To avoid the extra index I/O (machine time) that the query requires for a full-indexed evaluation, use the following SAS code:

```
PROC SQL;  
create table hr80spds  
as select  
  state,  
  age,  
  sex,  
  hour89,  
  industry,  
  occup  
from spdsцен.precs(wherenoindeх=(stser hour89))  
where occup='022'  
and state in('37','03','06','36')  
and industry='012'  
and hour89 > 40;  
quit;
```

Note: Specify index names in the WHERENOINDEX list, not the column names. The example excludes both the composite index for the STATE column STSER and the simple index HOUR89 from consideration by the WHINIT WHERE planner.

Options for Other Functions

BYSORT=

Perform an implicit automatic sort when SAS Scalable Performance Data (SPD) Server encounters a BY clause for a given table.

Syntax

```
BYSORT=YES | NO
```

Use the following arguments:

YES

sorts the data based on the BY columns and returns the sorted data to the SAS client. This powerful

capability means the user does not have to sort data using a PROC SORT statement before using a BY clause.

NO

does not sort the data based on the BY columns. This may be desirable if a DATA step BY clause has a GROUPFORMAT option or if a PROC step reports grouped and formatted data.

Description

The default is YES. The NO argument means the table must have been previously sorted by the requested BY columns. The NO argument allows grouped data to maintain their precise order in the table. A YES argument groups the data correctly but possibly in a different order from the order in the table.

Example 1 - Group Formatting with BYSORT=

```
libname sport sasspds 'mylib'
  host='samson'
  user='user19'
  passwd='dummy2';

PROC FORMAT;
  value dollars
    0-99.99="low"
    100-199.99="medium"
    200-1000="high";
run;

data sport.racquets;
  input
    raqname $20.
    @22 weight
    @28 balance $2.
    @32 flex
    @36 gripsize
    @42 string $3.
    @47 price
    @55 instock;

  datalines;
Solo Junior          10.1  N   2  3.75  syn   50.00  6
Solo Lobber          11.3  N  10  5.5   syn  160.00  1
Solo Queensize      10.9  HH   6  5.0   syn  130.00  3
Solo Kingsize       13.1  HH   5  5.6   syn  140.00  3
;

PROC PRINT data=sport.racquets (bysort=yes);
  var raqname instock;
  by price;
  format price dollars.;
  title 'Solo Brand Racquets by Price Level';
run;
```

Output 4.1 Report Output with BYSORT=

Solo Brand Racquets by Price Level		
----- Price=low -----		
OBS	RAQNAME	INSTOCK
1	Solo Junior	6
----- Price=medium -----		
OBS	RAQNAME	INSTOCK
3	Solo Queensize	3
4	Solo Kingsize	3
2	Solo Lobber	1

Example 2 - Group Formatting without BYSORT=

```
PROC PRINT data=sport.racquets (bysort=no);
  var raqname instock;
  by price;
  format price dollars.;
  title 'Solo Brand Racquets by Price Level';
run;
```

Output 4. 2 Report Output without BYSORT=

Solo Brand Racquets by Price Level		
----- Price=low -----		
OBS	RAQNAME	INSTOCK
1	Solo Junior	6
----- Price=medium -----		
OBS	RAQNAME	INSTOCK
2	Solo Lobber	1
3	Solo Queensize	3
4	Solo Kingsize	3

ENDOBS=

Specifies the end row (observation) number in a user-defined range for the processing of a given table.

Syntax

ENDOBS=*n*

Use the following argument:

n

is the number of the end row.

Description

By default, SAS Scalable Performance Data (SPD) Server processes the entire table unless the user specifies a range of rows with the STARTOBS= and ENDOBS= options. If the STARTOBS= option is used without the ENDOBS= option, the implied value of ENDOBS= is the end of the table. When both options are used together, the value of ENDOBS= must be greater than STARTOBS=.

In contrast to the Base SAS software options FIRSTOBS= and OBS=, the STARTOBS= and ENDOBS= SAS Scalable Performance Data (SPD) Server options can be used for WHERE Clause processing in addition to table input operations.

Example

Print only rows 2-4 of the table INVENTORY.OLD_AUTOS.

```
libname inventory sasspds 'conversion_area'
  server=husky.5105
  user='siteusr1'
  prompt=yes;

data inventory.old_autos;
  input
    year $4.
    @6 manufacturer $12.
    model $10.
    body_style $5.
    engine_liters
    @39 transmission_type $1.
    @41 exterior_color $10.
    options $10.
    mileage conditon;

  datalines;

1971 Buick      Skylark   conv   5.8   A   yellow   00000001 143000 2
1982 Ford      Fiesta   hatch  1.2   M   silver   00000001 70000 3
1975 Lancia    Beta     2door  1.8   M   dk blue  00000010 80000 4
1966 Oldsmobile Toronado 2door  7.0   A   black    11000010 110000 3
1969 Ford      Mustang  sptfrf 7.1   M   red      00000111 125000 3
;

PROC PRINT data=inventory.old_autos (startobs=2 endobs=4);
run;
```

Output 4.3 Data in the Printed Output

1982	Ford	Fiesta	hatch	1.2	M	silver	00000001	70000	3
1975	Lancia	Beta	2door	1.3	M	dk blue	00000010	80000	4
1966	Oldsmobile	Toronado	2door	7.0	A	black	11000010	110000	3

STARTOBS=

Specifies the start row (observation) number in a user-defined range for the processing of a given table.

Syntax

```
STARTOBS=n
```

Use the following argument:

n
is the number of the start row.

Description

By default, SAS Scalable Performance Data (SPD) Server processes the entire table unless the user specifies a range of rows with the STARTOBS= and ENDOBS= options. If the ENDOBS= option is used without the STARTOBS= option, the implied value of STARTOBS= is 1. When both options are used together, the value of STARTOBS= must be less than ENDOBS=.

In contrast to the Base SAS software options FIRSTOBS= and OBS=, the STARTOBS= and ENDOBS= SAS Scalable Performance Data (SPD) Server options can be used for WHERE Clause processing in addition to table input operations.

Example

Print only rows 2-4 of the table INVENTORY.OLD_AUTOS.

```
libname inventory sasspds 'conversion_area'
  server=husky.5105
  user='siteusr1'
  prompt=yes;
```

```
data inventory.old_autos;
  input
    year $4.
    @6 manufacturer $12.
    model $10.
    body_style $5.
    engine_liters
    @39 transmission_type $1.
    @41 exterior_color $10.
    options $10.
    mileage conditon;
```

```
datalines;
```

```
1971 Buick      Skylark  conv  5.8  A  yellow  00000001 143000 2
```

```

1982 Ford      Fiesta   hatch 1.2 M silver  00000001  70000 3
1975 Lancia   Beta    2door 1.8 M dk blue  00000010  80000 4
1966 Oldsmobile Toronado 2door 7.0 A black   11000010 110000 3
1969 Ford     Mustang sptrf 7.1 M red    00000111 125000 3
;

proc print data=inventory.old_autos (startobs=2 endobs=4);
run;

```

UNIQUESAVE=

Specifies to save rows with non-unique key values (the rejected rows) to a separate table when appending data to tables with unique indexes.

Syntax

```
UNIQUESAVE=YES|NO|REP
```

Default

```
NO
```

Complements the Table Option

```
SYNCADD=
```

Used in Conjunction with Macro Variable

```
SPDSUSDS=
```

Corresponding Macro Variable:

```
SPDSUSAV=
```

Use the following arguments:

YES

writes rejected rows to a separate, system-created table file which can be accessed by a reference to the macro variable SPDSUSDS=.

NO

does not write rejected rows to a separate table, that is, ignores non-unique key values.

REP

when updating a master table from a transaction table, where the two tables share identical variable structures, the UNIQUESAVE=REP option replaces the row updated row in the master table instead of appending a row to the master table. The REP option only functions in the presence of a /UNIQUE index on the MASTER table. Otherwise the REP setting is ignored..

Description

SYNCADD= is defaulted to NO. When NO, table appends are 'pipelined', meaning that the server data is sent in a stream a block at a time (see table option [NETPACKSIZE=](#)). While pipelining is faster than a synchronous append, SAS reports the results of the append operation differently for these two modes.

When applying only a single row (SYNCADD=NO), SAS returns a status code for each ADD operation. The application can determine the next action based upon the status value. If a row is rejected due to containing a non-unique value for a unique index, the user receives a status message. In contrast, when data is pipelined (SYNCADD=YES), SAS returns a status code only after **all** the rows are applied to a table. As a consequence, the user does not know which rows have been rejected.

To enjoy the performance of data pipelining but still retain the rejected rows, use the UNIQUESAVE= option. When set to YES, SAS Scalable Performance Data (SPD) Server will save any rows that are rejected to a hidden SAS table.

When using this option, SAS returns the name of the hidden table containing the rejected rows in the macro variable SPDSUSDS. If you want to report the contents of the table, reference [SPDSUSDS=](#).

Note: If SYNCADD= YES is set, data pipelining is overridden and the data is processed synchronously. In this situation, the UNIQUESAVE= option is not relevant and, if set, is ignored.

Example 1

We want to append two tables, NAMES2 and NAMES3, which contain employees' names, to the NAMES1 table. Before performing our append, we create an index on the NAME column in NAMES1, declaring the index unique.

Specify for SAS Scalable Performance Data (SPD) Server, during the append operation, to store rows found with duplicate employee names to a separate table file generated by the macro variable SPDSUSDS=.

Use a %PUT statement to display the table name for SPDSUSDS=. Then request a printout of the duplicate rows to review later.

```
data employee.names1;
input name $ exten;
datalines;
Jill 4344
Jack 5589
Jim 8888
Sam 3334
;
run;
```

```
data employee.names2;
input name $ exten;
datalines;
Jack 4443
Ann 8438
Sam 3334
Susan 5321
Donna 3332
;
run;
```

```
data employee.names3;
input name $ exten;
datalines;
Donna 3332
Jerry 3268
Mike 2213
;
run;
```

```
PROC DATASETS lib=employee nolist;
  modify names1;
```

```

    index create name/unique;
quit;

PROC APPEND data=employee.names2
            out=employee.names1(uniquesave=yes); run;

title 'The NAMES1 table with unique names
      from NAMES2';

PROC PRINT data=employee.names1;
run;

%put Set the macro variable spdsusds to &spdsusds;

title 'Duplicate (non-unique) name rows found in
      NAMES2';

PROC PRINT data=&spdsusds;
run;

PROC APPEND data=employee.names3
            out=employee.names1(uniquesave=yes);
run;

```

The SAS log provides the messages:

```

WARNING: Duplicate values not allowed on index NAME for
        file EMPLOYEE.NAMES1. (Occurred 2 times.)
NOTE: Duplicate records have been stored in file
      EMPLOYEE._30E3FD5.

```

And, an extract from our PROC PRINT shows:

The NAMES1 table with unique names from NAMES2

OBS	NAME	EXTENS
1	Jill	4344
2	Jack	5589
3	Jim	8888
4	Sam	3334
5	Ann	8438
6	Susan	5321
7	Donna	3332

Duplicate (non-unique) name rows found in NAMES2

OBS	NAME	EXTEN	XXX00000
1	Jack	4443	NAME
2	Sam	3334	NAME

Example 2

Use the UNiquesave=REP option to perform an update / append case using PROC APPEND instead of a DATA step:

```

/* A MASTER table to update. ID */
/* will get a UNIQUE index      */

DATA SPDS.MASTER;
  INPUT ID VALUE $;
  CARDS;
    1 one
    2 two
    3 three
  ;

PROC DATASETS LIB=SPDS;
  MODIFY MASTER;
  INDEX CREATE ID/UNIQUE;
QUIT;

/* A transaction table TRANS to use to */
/* drive update/appends to MASTER      */

DATA SPDS.TRANS;
  INPUT ID VALUE $;
  1 ONE
  3 THREE
  4 FOUR
  4 FOUR*
  ;

/* Use of UNIQUESAVE=REP to update/append */
/* TRANS rows to MASTER based on whether */
/* TRANS records have an ID column that */
/* matches an existing row from the MASTER */
/* table. Update MASTER rows with a match, */
/* otherwise append TRANS row to MASTER */

PROC APPEND DATA=SPDS.TRANS
  OUT=SPDS.MASTER(UNIQUESAVE=REP);
run;

```

Output of the resulting MASTER table would look like:

Obs	ID	VALUE
1	1	ONE
2	2	two
3	2	THREE
4	4	FOUR*

VERBOSE=

Provides details of all indexes and ACL information associated with a SAS Scalable Performance Data (SPD) Server table.

Syntax

VERBOSE= YES | NO

Use the following arguments:

YES

requests detail information for the indexes, ACLs, and other SAS Scalable Performance Data (SPD) Server table values. This argument must be used with the CONTENTS procedure.

NO

suppresses detail information for the indexes, ACLs, and other SAS Scalable Performance Data (SPD) Server table values. This is the default.

Example

Request details of all the indexes for the table TEMP1 in the domain SPDS44.

```
PROC CONTENTS data=SPDS44 (verbose=yes);  
run;
```

The CONTENTS Procedure

Data Set Name	SPDS44.TEMP1	
Observations	1000	
Member Type	DATA	
Variables	2	
Engine	SASSPDS	
Indexes	2	
Created	Tuesday, May 10, 2005 10:00:02 AM	Observation
Length	16	
Last Modified	Tuesday, May 10, 2005 11:01:36 AM	Deleted
Observations	0	
Protection		
Compressed	NO	
Data Set Type		
Sorted	NO	
Label		
Data Representation	Default	
Encoding	Default	

Engine / Host Dependent Information

Blocking Factor (obs/block)	2047
ACL Entry	NO
ACL User Access(R,W,A,C)	(Y,Y,Y,Y)
ACL User Name	ANONYMOU
ACL Owner Name	ANONYMOU
Data Set is Ranged	NO
Alphabetic List of Index Info	.
Bitmap index (No Global Index)	i
Keyvalue (Min)	1
Keyvalue (Max)	100
# of Discrete values	100
Bitmap index (No Global Index)	j
Keyvalue (Min)	1
Keyvalue (Max)	10
# of Discrete values	10
Data Partsize	16777216

Alphabetic List of Variables and Attributes

*	Variable	Type	Len
1	i	Num	8
2	j	Num	9

Alphabetic List of Indexes and Attributes

*	Index	# of Unique Values
1	i	100
2	j	10

Options for Security

ENCRYPT=

Encrypts SAS Scalable Performance Data (SPD) Server tables on disk. Encryption is a security mechanism that protects table contents from users who have system access to raw SAS Scalable Performance Data (SPD) Server tables. Access to tables is normally controlled by SAS Scalable Performance Data (SPD) Server ACLs. The SAS Scalable Performance Data (SPD) Server Administrator's Guide contains detailed information about using SAS Scalable Performance Data (SPD) Server ACLs to control access to tables.

When the ENCRYPT= option setting is set to YES, SAS Scalable Performance Data (SPD) Server encrypts newly created tables by blocks. To control the amount of encryption per block, use the table option [IOBLOCKSIZE=](#). The IOBLOCKSIZE= option specifies the number of rows to be encrypted in each block.

Syntax

ENCRYPT= YES | NO

Arguments

YES

encrypts the data set. The encryption method uses passwords. At a minimum, you must specify the READ= or the PW= data set option at the same time that you specify an ENCRYPT=YES option setting.

NO

no table encryption is performed. NO is the default setting for the ENCRYPT= option.

Usage Notes

1. Depending on your query patterns, increasing or decreasing the block size can affect performance.
2. SAS Scalable Performance Data (SPD) Server does *not* encrypt table indexes or metadata. Only table row data are encrypted.
3. To encrypt SPD tables with pass-through SQL, use only the READ= or PW= table option. With pass-through SQL, ENCRYPT=YES is implied with these options.
4. To access an encrypted table, the user must have appropriate ACL permissions to the table and must provide the encryption key via the READ= or PW= table option.
5. Encrypting a SAS Scalable Performance Data (SPD) Server table provides security from users that have system access to dump raw SAS Scalable Performance Data (SPD) Server tables. The section on Security in

the SAS Scalable Performance Data (SPD) Server Administrator's Guide contains more information about how to control system access to SAS Scalable Performance Data (SPD) Server tables.

Copyright © 2008 by SAS Institute Inc., Cary, NC, USA. All rights reserved.

SAS Scalable Performance Data (SPD) Server Formats and Informats

- [Introduction](#)
 - [Formats](#)
 - [Formats Example](#)
 - [User-Defined Formats Example](#)
 - [Informats](#)
-

[Introduction](#)

SAS Scalable Performance Data (SPD) Server supports some of the more commonly used SAS format and informats. Use these in your SQL pass-through code when you want SAS Scalable Performance Data (SPD) Server to associate a data set variable with a specific format.

A general reminder about formats: A format is applied to data set variables as it is written out. Informats are applied as the data set variable is being read. For more information on the formats and informats listed here see Chapters 13 and 14 in the *SAS Language Reference*.

[Formats](#)

- \$ -- Writes standard character data
- \$BINARY -- Converts character values to binary representation
- \$CHAR -- Writes standard character data
- \$HEX -- Converts character values to hexadecimal representation
- \$OCTAL -- Converts character values to octal representation
- \$QUOTE -- Converts character values to quoted strings
- \$VARYING -- Writes varying length values
- BEST -- SAS Scalable Performance Data (SPD) Server system chooses best notation
- BINARY -- Converts numeric values to binary representation
- COMMA -- Writes numeric values with commas and decimal points
- COMMAX -- Writes numeric values with commas and decimal points (European style)
- DATE -- Writes date values (ddmmyy)
- DATETIME -- Writes date time values (ddmmyy:hh:mm:ss.ss)
- DAY -- Writes day of month
- DDMYY -- Writes date values (ddmmyy)
- DOLLAR -- Writes numeric values with dollar signs, commas, and decimal points
- DOLLARX -- Writes numeric values with dollar signs, commas, and decimal points (European style)
- DOWNAME -- Writes name of day of the week
- E -- Writes scientific notation
- F -- Writes scientific notation
- FRACT -- Converts values to fractions
- HEX -- Converts real binary (floating-point) numbers to hexadecimal representation
- HHMM -- Writes hours and minutes
- HOUR -- Writes hours and decimal fractions of hours

- IB -- Writes integer binary values
- MMDDYY -- Writes date values (mmddyy)
- MMSS -- Writes minutes and seconds
- MMY Y -- Writes month and year, separated by a 'M'
- MONNAME -- Writes name of month
- MONTH -- Writes month of year
- MONYY -- Writes month and year
- NEG PAREN -- Displays negative values in parentheses
- OCTAL -- Converts numeric values to octal representation
- PD -- Writes packed decimal data
- PERCENT -- Prints numbers as percentages
- PIB -- Writes positive integer binary values
- QTR -- Writes quarter of year
- RB -- Writes real binary (floating-point) data
- SSN -- Writes social security numbers
- TIME -- Writes hours, minutes, and seconds
- TOD -- Writes the time portion of datetime values
- *w.d* -- Writes standard numeric data
- WEEKDATE -- Writes day of week and date (day-of-week, month-name dd, yy)
- WEEKDATX -- Writes day of week and date (day-of-week, dd month-name yy)
- WEEKDAY -- Writes day of week
- WORDDATE -- Writes date with name of month, day, and year (month-name dd, yyyy)
- WORDDATX -- Writes date with day, name of month, and year (dd month-name yyyy)
- WORDF -- Converts numeric values to words
- WORDS -- Converts numeric values to words (fractions as words)
- YEAR -- Writes year part of date value
- YYMM -- Write year and month, separated by a 'M'
- YYMMDD -- Writes day values (yymmdd)
- YYMON -- Writes year and month abbreviation
- YYQ -- Writes year and quarter, separated by a 'Q'
- Z -- Writes leading 0s
- ZD -- Writes data in zoned decimal format

Note: Formats which begin with a '\$' sign are character formats. Otherwise the format accepts numeric values.

Formats Example:

Use the *dollar.* format to convert numeric sales figures into dollar values. Suppose you have a SAS Scalable Performance Data (SPD) Server data set (sales) with a single numeric variable (salesite) representing the total sales for a given site. Using SQL pass-through, create a new data set containing the sales in dollar format.

```

PROC SQL;
connect to sasspds
  (dbq='tmp'
   user='anonymous'
   host='localhost'
   serv='5127');

execute(create table money
  as select salesite
  format=dollar.
  from sales)

```

```
by sasspds;  
  
quit;
```

User-Defined Formats Example

This example is a sample test job that validates its own configuration to utilize user-defined formats. When properly configured, user-defined formats will allow columns to be formatted using Parallel Group-By statements and a WHERE-clause that uses a format to subset data to the server.

The example provides sample spdserv.parm and libnames.parm file examples as well as code examples which follow the two sample SAS Scalable Performance Data (SPD) Server configuration files.

This example is a sample test job that checks the usage of user-defined formats. When correctly set up, user-defined formats will allow formatting of columns in parallel group-by and permits usage of a WHERE-clause that uses a format to subset data.

SAS Scalable Performance Data (SPD) Server spdserv.parm file used in the example:

```
SORTSIZE=8M;  
INDEX_SORTSIZE=8M;  
BINBUFSIZE=32K;  
INDEX_MAXMEMORY=8M;  
NOCOREFILE;  
SEQIOBUFMIN=64K;  
RANIOBUFMIN=4K;  
NOALLOWMMAP;  
MAXWHTHREADS=16;  
WHERECOSTING;  
RANDOMPLACEDPF;  
FMTDOMAIN=FORMATS;  
FMTNAMENODE=d8488 ;  
FMTNAMEPORT=5200;
```

SAS Scalable Performance Data (SPD) Server libnames.parm file used in the example:

```
libname=tmp pathname=c:\temp;  
libname=formats pathname=c:\data\formats;
```

SAS Scalable Performance Data (SPD) Server example code:

```
%let domain=tmp;  
%let host=d8488;  
%let serv=5200;  
  
/* locking=YES must be specified when using */  
/* options fmtsearch=(formats); */  
  
libname formats sasspds 'formats'  
    host=&host  
    serv=&serv  
    user='anonymous'
```

```

locking=YES;

libname &domain sasspds "&domain"
    host="&host"
    serv="&serv"
    user='anonymous'
    IP=YES;

options fmtsearch=(formats);

PROC DATASETS nolist
    lib=formats
    memtype=catalog;

delete formats;

quit ;

/* To create user defined formats, they must be */
/* loaded from the same platform where they are */
/* going to be stored. You cannot use Windows */
/* path specifications to load formats on UNIX */
/* platforms. */

/* Add formats to format domain */

PROC FORMAT lib=formats;
value AGEGRP
    0-13='Child'
    14-17='Adolescent'
    18-64='Adult'
    65-HIGH='Pensioner';

value $GENDER
    'F' = 'Female'
    'M' = 'Male';

run ;

/* Create a test table with a column that uses */
/* AGEGRP format */

data &domain..fmttest;
format age AGEGRP. GENDER $GENDER. id z5.;
    length GENDER $1;
do id=1 to 100;
    if mod (id,2) = 0
        then GENDER = 'F';
        else GENDER = 'M';
    age=int(ranuni(0)*100);
    income=age*int(ranuni(0)*1000);
output;

end;

```

```

run;

/* Use the Parallel Group-By feature with the      */
/* fmtgrp sel option. This groups the data based */
/* on the output format specified in the table. */
/* This will be executed in parallel.           */

PROC SQL;
connect to sasspds
  (dbq="&domain"
   serv="&serv"
   host="&host"
   user="anonymous");

/* Explicitly set the fmtgrp sel option          */

execute(reset fmtgrp sel)
by sasspds;

title 'Simple Fmtgrp sel Example';

select *
from connection to sasspds
  (select age, count(*) as count
   from fmtttest group by age);

disconnect from sasspds;

quit;

PROC SQL;
connect to sasspds
  (dbq="&domain"
   serv="&serv"
   host="&host"
   user="anonymous");

/* Explicitly set the fmtgrp sel option          */

execute(reset fmtgrp sel)
by sasspds;

title 'Format Both Columns Group Select Example';

select *
from connection to sasspds
  (select GENDER format=$GENDER.,
         AGE format=AGEGRP.,
         count(*) as count
   from fmtttest
   formatted group by GENDER, AGE);

disconnect from sasspds;

quit;

```

```

PROC SQL;
connect to sasspds
  (dbq="&domain"
   serv="&serv"
   host="&host"
   user="anonymous");

/* Explicitly set the fmtgrpsel option */

execute(reset fmtgrpsel)
by sasspds;

title1 'To use Format on Only One Column With Group Select';
title2 'Override Column Format With a Standard Format';

select *
from connection to sasspds
  (select GENDER format=$1.,
        AGE format=AGEGRP.,
        count(*) as count
   from fmttest
   formatted group by GENDER, AGE);

disconnect from sasspds;

quit;

/* A WHERE-clause that uses a format to subset */
/* data is pushed to the server. If it is not */
/* pushed to the server, the following warning */
/* message will be written to the SAS log: */
/* WARNING: Server is unable to execute the */
/* where clause. */

data temp;
set &domain..fmttest;
  where put
    (AGE,AGEGRP.) = 'Child';
run;

```

Informats

- \$ -- Reads standard character data
- \$BINARY -- Converts binary values to character values
- \$CB -- Reads standard character data from column-binary files
- \$CHAR -- Reads character data with blanks
- \$HEX -- Converts hexadecimal data to character data
- \$OCTAL -- Converts octal data to character data
- \$PHEX -- Converts packed hexadecimal data to character data
- \$QUOTE -- Converts quoted strings to character data
- \$SASNAME --
- \$VARYING -- Reads varying length values
- BEST -- SAS Scalable Performance Data (SPD) Server system chooses best notation

- **BINARY** -- Converts positive binary values to integers
- **BITS** -- Extract bits
- **COMMA** -- Removes embedded characters (e.g. \$,.)
- **COMMAX** -- Removes embedded characters (e.g. \$,.) European style
- **D** -- Reads scientific notation
- **DATE** -- Reads date values (ddmmyy)
- **DATETIME** -- Reads datetime values (ddmmyy hh:mm:ss.ss)
- **DDMMYY** -- Reads date values (ddmmyy)
- **DOLLAR** -- Reads numeric values with dollar signs, commas, and decimal points
- **DOLLARX** -- Reads numeric values with dollar signs, commas, and decimal points (European style)
- **E** -- Reads scientific notation
- **F** -- Reads scientific notation
- **HEX** -- Converts hexadecimal positive binary values to fixed- or floating-point values
- **IB** -- Reads integer binary (fixed-point) values
- **JULIAN** -- Reads Julian dates (yyddd or yyyyddd)
- **MMDDYY** -- Reads date values (mmddy)
- **MONYY** -- Reads month and year date values (mmyy)
- **MSEC** -- Reads TIME MIC values
- **OCTAL** -- Converts octal values to integers
- **PD** -- Reads packed decimal data
- **PDTIME** -- Reads packed decimal time of SMF and RMF records
- **PERCENT** -- Converts percentages into numeric values
- **PIB** -- Reads positive integer binary (fixed-point) values
- **PK** -- Reads unsigned packed decimal data
- **PUNCH** -- Reads whether a row of column-binary data is punched
- **RMFSTAMP** -- Reads time and date fields of RMF records
- **ROW** -- Reads a column-binary field down a card column
- **SMFSTAMP** -- Reads time-date values of SMF records
- **TIME** -- Reads hours, minutes, and seconds (hh:mm:ss.ss)
- **TODSTAMP** -- Reads 8-byte time-of-day stamp
- **TU** -- Reads timer units
- **YYMMDD** -- Reads date values (yymmdd)
- **YYQ** -- Reads quarters of the year

Note: Informats which begin with a '\$' sign are character informats. Otherwise the informat accepts numeric values.

The SQL procedure itself does not use the **INFORMAT=** modifier: it stores informats in its table definitions so that other procedures and the data step can use the information. SAS Scalable Performance Data (SPD) Server informats are provided now to allow for forward compatibility with future development.

Copyright © 2008 by SAS Institute Inc., Cary, NC, USA. All rights reserved.

SAS Scalable Performance Data (SPD) Server NLS Support

Contents

- [Overview of NLS](#)
- [Character Encoding Overview](#)
 - [What is Character Encoding?](#)
 - [Common Encodings](#)
- [Moving Data across Environments with Different Encodings](#)
 - [Transcoding](#)
 - [How Base SAS Transcodes Data](#)
 - [Base SAS Encoding Behavior](#)
 - [SAS 9 Output Processing](#)
 - [SAS 9 Input Processing](#)
 - [Reading and Writing External Files](#)
- [Setting the Encoding for Base SAS Sessions](#)
- [Changing the Encoding for Base SAS Sessions](#)
- [NLS Support in SAS Scalable Performance Data \(SPD\) Server](#)
- [SAS Scalable Performance Data \(SPD\) Server NLS Limitations](#)
 - [Affected Data](#)
 - [Pass-Through SQL](#)
 - [Case Folding and Sort Sequences](#)
 - [Indexes and Ordering](#)
 - [Date and Time Representations](#)
 - [Suppressing Transcoding](#)
 - [LIBNAME Option Restrictions:](#)

[Overview of NLS](#)

NLS, or National Language Support, deals both with *Internationalization* and *Localization* of SAS software. *Internationalization* is the process of designing an application so that it can be adapted to different languages and regions, without requiring engineering changes. Often the term internationalization is abbreviated as **i18n**, because there are 18 letters between the first "i" and the last "n." *Localization* is the process of adapting software for a particular region or language by adding locale-specific components and translating text. The term localization is often abbreviated as **L10n**, because there are 10 letters between the "L" and the "n." Translation of user interface, messages, and documentation is a large part (but not all) of localization. Localizers also verify that the formatting of dates, numbers, currencies etc. conforms to local requirements.

SAS 9 contains built-in support for NLS character set encoding and locale choices. Users access the NLS encoding and locale choices through various SAS, LIBNAME, and data set options. SAS Scalable Performance Data (SPD) Server and SAS together offer basic levels of NLS support. This document describes the basic entities of NLS support and how they are implemented in SAS Scalable Performance Data (SPD) Server

Character Encoding Overview

All input to a computer is represented internally as numbers. The computer assigns a number to each character -- technically, the number is a binary number (base 2 numbering system, consisting of 0s and 1s).

Because most of us don't want to think in binary numbers, computers provide hexadecimal (base 16 numbering system) representation as a shorthand for binary representation. For example, for the decimal number 167, it's easier to understand the hexadecimal number A7 than the equivalent binary number 10100111. Therefore, you can think of the computer's internal numeric representation of all data as a hexadecimal number.

What is Character Encoding?

All data that is stored, transmitted, or processed by a computer is in an encoding. An *encoding* maps each character to a unique numeric representation. For example:

1. You press a key on a keyboard, like the uppercase letter A.
2. The computer assigns the internal numeric representation, that is, a unique hexadecimal number.
3. To display or print the character, the computer uses the font (graphical representation) that matches the numeric representation, that is, the uppercase letter A.

To assign the numeric representation to a character, an encoding uses a *code page*, which is an ordered set of characters in which a numeric index (code point value) is associated with each character. The position of a character on the code page determines its two-digit hexadecimal number. The first digit of the hexadecimal number is determined by the column, and the second digit by the row. For example, the following is the code page for the Windows Latin1 encoding. The numeric representation for the uppercase A is the hexadecimal number 41, and the numeric representation for the equal sign (=) is the hexadecimal number 3D.

HEX DIGITS 1ST → 2ND ↓	0-	1-	2-	3-	4-	5-	6-	7-	8-	9-	A-	B-	C-	D-	E-	F-
-0			0	@	P	'	p	€			ø	°	À	Ð	á	ð
-1			1	A	Q	a	q		*	ı	±	Á	Ñ	á	ñ	
-2			"	2	B	R	b	r	,	'	€	²	Â	Ò	â	ò
-3			#	3	C	S	c	s	f	"	£	³	Ã	Ó	ã	ó
-4			\$	4	D	T	d	t	,	"	¤	'	Ä	Ö	ä	ö
-5			%	5	E	U	e	u	...	•	¥	µ	Å	Ö	å	ö
-6			&	6	F	V	f	v	†	-		¶	Æ	Ö	æ	ö
-7			'	7	G	W	g	w	‡	—	§	•	Ç	×	ç	÷
-8			(8	H	X	h	x	^	~	"	,	È	Ø	è	ø
-9)	9	I	Y	i	y	‰	™	©	ı	É	Ù	é	ù
-A			*	:	J	Z	j	z	Š	š	•	•	Ê	Ú	ê	ú
-B			+	;	K	[k	{	<	>	«	»	Ë	Û	ë	û
-C			,	<	L	\	l		œ	œ	¬	¼	Ì	Ü	ì	ü
-D			-	=	M]	m	}			½	½	Í	Ý	í	ý
-E			.	>	N	^	n	~	Ž	ž	®	¾	Î	Þ	î	þ
-F			/	?	O	_	o				™	™	Ï	ß	ï	ÿ

Encoding is the combination of a character set with an encoding method:

- A *character set* is the repertoire of characters and symbols that are used by a language or group of languages. A character set includes national characters (which are characters specific to a particular nation or group of nations), special characters (such as punctuation marks), the unaccented Latin characters A-Z, the digits 0-9, and control characters that are needed by the computer.
- An *encoding method* is the set of rules that are used to assign the numbers to the set of characters that are in an encoding. These rules govern such things as the size of the encoding (number of bits used to store the numeric representation of the character) and the ranges in the code page where characters are allowed to appear.

When the rules of the encoding method are followed, and numbers are assigned to the characters,

the result is called an encoding.

An individual character can have different positions in code pages for different encodings, which result in different hexadecimal numbers. For example, the position of the uppercase letter A in the Wlatin1 code page (shown above) results in the hexadecimal number 41, while in the following Danish EBCDIC code page, the position of the uppercase letter A results in the hexadecimal number C1.

HEX DIGITS 1ST → 2ND ↓	4-	5-	6-	7-	8-	9-	A-	B-	C-	D-	E-	F-
-0	SP BF010000	& SM090000	- SP100000	 SM500000	@ BM090000	° BM100000	μ BM170000	¢ BC040000	æ LA610000	ä LA270000	\ SM070000	0 ND100000
-1	QSP SP900000	ó LE110000	/ SP120000	É LE120000	a LA010000	j LJ010000	ū LJ170000	£ BC090000	A LA020000	J LJ020000	+ SA060000	1 ND010000
-2	â LA150000	ë LE150000	Â LA160000	Ê LE160000	b LB010000	k LX010000	s LB010000	¥ SD060000	B LB020000	K LX020000	S LB020000	2 ND020000
-3	ã LA170000	ë LE170000	Ä LA180000	Ë LE180000	c LD010000	l LL010000	t LT010000	· SD090000	C LC020000	L LL020000	T LT020000	3 ND030000
-4	á LA130000	è LE130000	À LA140000	È LE140000	d LD010000	m LM010000	u LU010000	© SM520000	D LD020000	M LM020000	U LU020000	4 ND040000
-5	á LA110000	í LJ110000	Á LA120000	Í LJ120000	e LE010000	n LN010000	v LV010000	§ SM240000	E LE020000	N LN020000	V LV020000	5 ND050000
-6	â LA190000	î LJ150000	Ã LA200000	Ï LJ160000	f LF010000	o LO010000	w LW010000	¶ SM260000	F LF020000	O LO020000	W LW020000	6 ND060000
-7	} SM140000	ÿ LJ170000	Š SC090000	Ï LJ180000	g LG010000	p LP010000	x LX010000	¼ NF040000	G LG020000	P LP020000	X LX020000	7 ND070000
-8	ç LC010000	ì LJ190000	Ç LC020000	Ì LJ140000	h LH010000	q LQ010000	y LY010000	½ NF010000	H LH020000	Q LQ020000	Y LY020000	8 ND080000
-9	ñ LN100000	ß LS010000	Ñ LN020000	´ SD130000	ï LI010000	r LR010000	z LZ010000	¾ NF050000	I LI020000	R LR020000	Z LZ020000	9 ND090000
-A	# SM010000	€ SC230000	ø LC010000	: SP130000	« SP170000	ª SM210000	¡ SP030000	¬ SM660000	Š SP200000	1 ND010000	2 ND020000	3 ND030000
-B	· SP110000	À LA280000	· SP080000	Æ LA520000	» SP180000	º SM200000	¿ SP160000	¡ SM130000	ó LC150000	ô LJ150000	õ LC160000	û LJ160000
-C	< SA090000	* SM040000	% SM020000	Ø LC020000	ø LD030000	{ SM110000	Ð LD620000	- SM150000	ö LC170000	~ SD100000	Ö LD180000	Û LJ180000
-D	(SP080000) SP070000	_ SP090000	´ SP060000	ý LY110000	· SD470000	Ý LY120000	" SD170000	ö LD190000	ù LJ130000	ò LD140000	Û LJ140000
-E	+ SA010000	; SP140000	> SA050000	= SA040000	þ LT690000	[SM060000	Þ LT640000	' SD110000	ó LC110000	ú LJ110000	Ó LD120000	Ú LJ120000
-F	¡ SP020000	^ SD150000	? SP150000	" SP040000	± SA020000] SM060000	® SM590000	x SA070000	ö LC190000	ÿ LY170000	Ö LD200000	¸ BC010000

Common Encodings

There are many encodings that address the requirements of different languages. Very few languages use only the 26 characters A through Z of the Latin alphabet. In addition, there are

different encodings to address different operating system standards.

An encoding that represents each character in one byte is a single-byte character set (SBCS). A single-byte character set can be either 7 bits (providing up to 128 characters) or 8 bits (providing up to 256 characters). An example of an 8-bit SBCS is the Latin1 encoding (represents the characters of Western Europe). (Note that the term *octet*, for the international community, is an 8-bit byte. Since a byte is not 8 bits in all computer systems, octet provides an unambiguous term.)

A multiple-byte character set (MBCS) is a mixed-width encoding in which some characters consist of more than one byte. For example, the Japanese, Korean, Simplified Chinese, and Traditional Chinese are MBCS encodings. A double-byte character set (DBCS) is a specific type of a MBCS encoding that includes characters that consist of two bytes.

The following are common encodings:

ASCII (American Standard Code for Information Interchange)

is a 7-bit encoding for the United States that provides 128 character combinations. The encoding contains characters for uppercase and lowercase English, American English punctuation, base 10 numbers, and a few control characters. The set of 128 characters is the one common denominator that is contained in most encodings, excluding EBCDIC-based encodings. ASCII is used by personal computers.

ISO (International Organization for Standardization) 646 family

is a 7-bit encoding that is an international standard and provides 128 character combinations. The ISO 646 family of encodings is like ASCII except for 12 code points for national variants. The 12 national variants represent specific characters needed for a particular language.

EBCDIC (Extended Binary Coded Decimal Interchange Code) family

is an 8-bit encoding that provides 256 character combinations. There are multiple EBCDIC-based encodings. EBCDIC is used on IBM mainframes and most IBM midrange computers. EBCDIC follows ISO 646 conventions to facilitate translations between itself and 7-bit ASCII-based encodings. Characters A-Z and 0-9 are mapped to the same code points on all EBCDIC code pages, while the rest of the code points may be used for special characters and national characters, depending on the encoding.

ISO 8859 family and Windows family

is an 8-bit extension of ASCII that supports all of the ASCII code points and adds 12 more, providing 256 character combinations. Latin1, which is officially named ISO-8859-1, is the most frequently used member of the ISO 8859 family of encodings. In addition to the ASCII characters, Latin1 contains accented characters, other letters needed for languages of Western Europe, and some special characters.

Unicode

uses two bytes for each character rather than one and provides up to 65,536 character combinations. Unicode can handle the scripts of basically all of the world's languages. For

example, the Japanese language, which has thousands of characters, uses a 16-bit, multiple-byte character set. There are various forms of Unicode, including UTF-8, UTF-16, and UTF-32.

Moving Data across Environments with Different Encodings

Transcoding

Although it's easy to move data across environments that use the same encoding, it can be more difficult to move data across environments that use different encodings. When the encoding of a file is incompatible with the computer environment's encoding, transcoding occurs.

Transcoding is the process of mapping data from one encoding to another, for example, from an ASCII-based encoding to an EBCDIC-based encoding. Transcoding is not translating from one language to another; transcoding is *remapping of characters*.

For example, consider a file that was created on a UNIX platform that uses the Latin1 encoding, then moved to an IBM mainframe that uses the Danish EBCDIC encoding. When the file is processed on the IBM mainframe, the data is remapped from the Latin1 encoding to the Danish EBCDIC encoding. If the data contains a dollar sign (\$), the hexadecimal number is converted from 24 to 67.

Transcoding can occur in the following situations:

- when you move a SAS file from one platform to another and the file's encoding is incompatible with the current session encoding, for example, from an z/OS operating environment with an EBCDIC-based encoding to a Windows operating environment with an ASCII-based encoding.
- when you share data between two SAS sessions (like in a client/server environment) that have incompatible session encodings.
- when you read and write an external file.

How Base SAS Transcodes Data

Base SAS provides transcoding when you move data and applications from one environment to another. To transcode one encoding to another, SAS uses translation tables, like the one that maps Wlatin2 (Windows) to ISO Latin2 (UNIX).

For example, when you

- use the CPORT and CIMPORT procedures to create a transport file, SAS automatically uses translation tables to transcode one encoding to another and back again. First, the data is converted from the source encoding to transport format, then the data is converted from the transport format to the target encoding.
- process a SAS data set that has an encoding that is different from the current session encoding, SAS automatically uses CEDA (cross environment data access) software to transcode data. (CEDA is the same software in SAS that converts a SAS file to the correct data representation when you move a file from one platform to another.)

Base SAS Encoding Behavior

For base SAS files (not SAS Scalable Performance Data (SPD) Server), the encoding support depends on the version of SAS that created the file:

- Data sets created in SAS 9 automatically have an encoding attribute, which is stamped in the descriptor portion of the file.
- Data sets created in SAS 8 do not have an encoding value stamped on the file; they are assumed to be in the session encoding of the host environment.

The NLS features in SAS Scalable Performance Data (SPD) Server only support encoding from SAS 9.

SAS 9 Output Processing

For SAS 9 data sets (not SAS Scalable Performance Data (SPD) Server), encoding is determined as follows:

- For a new output file, the data is written to the file using the current session encoding.
- For a new output file that is created with the OUTREP= option, which specifies a data representation different from the current session, the data is written to the file using the default session encoding for the operating system that is based on the specified OUTREP= value.
- For output processing that replaces an existing file, the new file inherits the encoding of the existing file.

- For output processing that replaces an existing file that is from another platform or if the existing file has no encoding stamped on it, then the current session encoding is used.

SAS 9 Input Processing

For input (read) processing in SAS 9 (not SAS Scalable Performance Data (SPD) Server), encoding behavior is as follows:

- If the session encoding and the encoding that is stamped on the file are incompatible, the data is transcoded to the session encoding. For example, if the current session encoding is ASCII and the encoding that is stamped on the file is EBCDIC, SAS transcodes the data from EBCDIC to ASCII.
- If a file does not have an encoding stamped on it, SAS transcodes the data only if the file's data representation is different from the current session.

Reading and Writing External Files

SAS reads and writes external files using the current session encoding. SAS assumes that the external file is in the same encoding as the session encoding. For example, if you are creating a new SAS data set by reading an external file, SAS assumes that the file's encoding is the same as the session encoding. If it is not, the data could be written to the new SAS data set incorrectly.

Setting the Encoding for Base SAS Sessions

When SAS 9 is installed, the base SAS (not SAS Scalable Performance Data (SPD) Server) default encoding is host dependent and is determined by the default settings for several SAS system options. Here are three system options that you should be familiar with:

ENCODING=

establishes the session encoding, which is the encoding that SAS uses to process syntax, process SAS data sets, and read and write external files. The default value is host dependent; all are SBCS encodings:

Host	Value	Description
OpenVMS Alpha	Latin1	Western (ISO)

z/OS	OPEN_ED_1047	OpenEdition EBCDIC cp1047-Latin1
UNIX	Latin1	Western (ISO)
Windows	WLatin1	Western (Windows)

LOCALE=

specifies the locale of the SAS session. The locale reflects the local language, conventions, and culture for a particular geographical region. A locale's conventions may include the formatting of dates, times, and numbers, and printer preferences like paper size. Specifying a locale also automatically sets the default encoding that establishes the session encoding; a locale has a common encoding that is used most often for a particular operating environment. The default locale is English, and the common encodings for English are the defaults above for ENCODING=.

NONLSCOMPATMODE | NLSCOMPATMODE

provides national language compatibility for non-English data processing using native characters. For SAS 9, the default is NONLSCOMPATMODE, which provides consistency for running SAS on multiple systems. NONLSCOMPATMODE specifies that data is to be processed in the encoding that is set by the ENCODING= or LOCALE= system option.

[Changing the Encoding for Base SAS Sessions](#)

You can change the session encoding by using the LOCALE= system option, the ENCODING= system option, or both. Note that valid values for both options are host dependent.

Here's how you can set the base SAS (not SAS Scalable Performance Data (SPD) Server) session encoding when NONLSCOMPATMODE is specified:

- You can specify the LOCALE= system option in a configuration file, at SAS invocation, in an OPTIONS statement, or in the SAS System Options window. In SAS 9, several NLS-related system options are automatically set, based on the value of LOCALE=. Most customers will implicitly set encoding with the LOCALE= system option.
- You can specify the ENCODING= system option in a configuration file or at SAS invocation.
- Here is how LOCALE= and ENCODING= interact:
 - If a value is not specified for ENCODING= (that is, the installation default is set), then specifying a value for LOCALE= sets the encoding based on the LOCALE= value. In addition, values for the following system options are set based on the LOCALE= value: DFLANG=,

TRANTAB=, DATESTYLE=, and PAPERSIZE=.

- If a value is specified for ENCODING=, that value sets the session encoding and overrides LOCALE=.
 - If the value specified for LOCALE= is not compatible with the value specified for ENCODING=, then the value for LOCALE= is used. A warning message is provided if ENCODING= and LOCALE= conflict.
-
- If the DBCS system option is set, which specifies that SAS process DBCS encodings, the values for DBCSLANG= and DBCSTYPE= system options determine the session encoding and the locale. These options are used for Asian languages or for English with DBCS extensions.

Here is an example of implicitly setting the base SAS (not SAS Scalable Performance Data (SPD) Server) session encoding based on the specified locale when you invoke SAS:

```
sas9 -explorer -locale spanish
```

Here is an example of explicitly setting the base SAS (not SAS Scalable Performance Data (SPD) Server) session encoding with the OPTIONS statement:

```
options encoding=wlatin2;
```

Tip: Changing encoding for a SAS session does not affect SAS keywords, which remain in English, or SAS log output, which also remains in English.

[NLS Support in SAS Scalable Performance Data \(SPD\) Server](#)

SAS Scalable Performance Data (SPD) Server contains support for a subset of the SAS 9 NLS functions documented above. SAS Scalable Performance Data (SPD) Server utilizes encoding and locale currently only on SAS software.

In future releases of SAS Scalable Performance Data (SPD) Server, the locale identifier information will be used for locale-sensitive case folding and linguistic collation. Case-folding is defined as "a process applied to a sequence of characters, in which those identified as non-uppercase are replaced by their uppercase equivalents". Linguistic collation is performing linguistic sorts based on linguistic sort keys. However, those functions have yet to be implemented in SAS Scalable Performance Data (SPD) Server production code.

All tables that are produced by SAS Scalable Performance Data (SPD) Server and SAS inherit the SAS session's default encoding and locale settings. By default, SAS Scalable Performance Data (SPD) Server code expects new tables to follow the current SAS session's encoding and locale. Table updates that append rows or update existing rows will perform transcoding to ensure that appended and updated table rows match the existing table encoding.

Wire transfer is in the character set encoding of the SAS session for transfers to and from the SAS Scalable Performance Data (SPD) Server host, unless SAS Scalable Performance Data (SPD) Server transcoding has been disabled. SAS Scalable Performance Data (SPD) Server transcoding is enabled or disabled by inserting a [NO]NLSTRANS CODE statement in the SAS Scalable Performance Data (SPD) Server spdsserv.parm parameter file.

SAS Scalable Performance Data (SPD) Server NLS Limitations

Affected Data

SAS Scalable Performance Data (SPD) Server hosts are restricted in the way they handle NLS character strings. SAS Scalable Performance Data (SPD) Server hosts are restricted to data that is contained in character columns in data sets and some metadata structures. The NLS support for SAS Scalable Performance Data (SPD) Server is functional for only table labels and variable labels.

Column names, index names, table names, and catalog names are *not* supported in the SAS Scalable Performance Data (SPD) Server NLS support. Column names, index names, table names, and catalog names are still dependent on ASCII support. SAS Scalable Performance Data (SPD) Server SQL is subject to the NLS same restrictions.

Pass-Through SQL

SAS Scalable Performance Data (SPD) Server pass-through SQL does not support any NLS functions. Pass-through SQL operates in the encoding and locale of the SAS session that initiates the CONNECT to SASSPDS.

Case Folding and Sort Sequences

SAS Scalable Performance Data (SPD) Server NLS code supports very limited English Latin1 and Polish Latin2 case folding for SBCS encodings. UTF8 case folding is limited to the ASCII range of UTF8 encoding. NLS Sort sequences for SAS Scalable Performance Data (SPD) Server 4.4 are restricted to lexical sorts for all combinations. Linguistic sorting is a subject for future SAS Scalable Performance Data (SPD) Server releases.

Indexes and Ordering

Indexes in SAS Scalable Performance Data (SPD) Server are created in the table's encoding, and only support lexical ordering. If the client's encoding and locale settings match the SAS Scalable Performance Data (SPD) Server host table's encoding and locale settings, index use is unrestricted. Otherwise, index usage is restricted to certain predicates in WHERE clauses that can be safely interpreted according to the table's encoding and locale settings. When the client and host table encoding and locale settings differ, the EVAL2 strategy is used to filter predicates that require use of order.

Date and Time Representations

SAS Scalable Performance Data (SPD) Server server-side functions and formats that produce or accept textual date, time, and date/time representations are *not* locale-sensitive.

Suppressing Transcoding

You can suppress transcoding in the SAS Scalable Performance Data (SPD) Server environment by entering the following into the `spdsserv.parm` options:

```
NONLSTRANSICODE;
```

If you add the NONLSTRANSICODE option to your `spdsserv.parm` file, character transcoding between the SAS Scalable Performance Data (SPD) Server host and connected clients is disabled. Disabling character transcoding restricts the kinds of operations that the SAS Scalable Performance Data (SPD) Server host performs to operations it can safely perform, where host and client tables share the same encoding. Disabling SAS Scalable Performance Data (SPD) Server host transcoding assumes that the *client* will perform any needed transcoding on the data streams that it sends and receives to match the encoding of referenced tables. The SAS Scalable Performance Data (SPD) Server host setting for NONLSTRANSICODE does not perform any actions to deny client access to a host table that has mismatched encoding.

LIBNAME Option Restrictions:

The following options are not implemented in the SAS Scalable Performance Data (SPD) Server NLS functions:

The LIBNAME option

OUTENCODING=<client-server encoding>

is not supported and will produce a WARNING message if submitted to SASSPDS.

In addition, the related data set option

ENCODING=<client-server encoding>

is supported by the SAS LIBNAME engine for OUTPUT data sets only. Character data is assumed to be in the encoding of the session that initiates the CONNECT to SASSPDS and is normally stored using that encoding. ENCODING= will cause SAS Scalable Performance Data (SPD) Server to transcode from the SAS session encoding to the specified encoding for storing data. If you specify ENCODING= for a non OUTPUT data set open, and if the encoding value that you specify doesn't match the data set's encoding, the data set open will produce a warning:

ENCODING= specified on table open fails to match table encoding. Option ignored.

The LIBNAME option

TRUNCWARN=YES

Suppresses hard failure on NLS transcoding overflow and character mapping errors. When using the TRUNCWARN=YES LIBNAME option, data integrity may be compromised because significant characters can be lost in this configuration. The default setting is NO, which causes hard read/write stops when transcode overflow or mapping errors are encountered. When TRUNCWARN=YES, and an overflow or character mapping error occurs, a warning is posted to the SAS log at data set close time if overflow occurs, but the data overflow is lost.

Copyright © 2008 by SAS Institute Inc., Cary, NC, USA. All rights reserved.

SAS Scalable Performance Data (SPD) Server Frequently Asked Questions

Contents:

- [Does SAS Scalable Performance Data \(SPD\) Server support files greater than 2 Gigabytes?](#)
- [Can I create file systems greater than 2 Gigabytes?](#)
- [How do SAS Scalable Performance Data \(SPD\) Server client and server processes communicate?](#)
- [How do I know which ports must be surfaced through an Internet firewall?](#)
- [How does SAS Scalable Performance Data \(SPD\) Server interact with multi-homed hosts?](#)
- [Can I use standard UNIX backup procedures?](#)
- [What do I need to know about SAS Scalable Performance Data \(SPD\) Server installation? How long will it take?](#)
- [Is it necessary to run UNIX SAS Scalable Performance Data \(SPD\) Server as root?](#)
- [What is the Name Server and why do I need it?](#)
- [Does every SAS Scalable Performance Data \(SPD\) Server client need a UNIX ID or Windows Networking ID?](#)
- [Can a SAS Scalable Performance Data \(SPD\) Server host, Name Server and a SAS Scalable Performance Data \(SPD\) Server client all run on the same machine?](#)
- [Can I have multiple SAS Scalable Performance Data \(SPD\) Server hosts on the same machine?](#)
- [How do I create LIBNAME domains?](#)
- [How do I specify a LIBNAME domain in SAS?](#)
- [Is there anything else I have to change to run my existing SAS applications?](#)
- [How can I get existing data loaded into a SAS Scalable Performance Data \(SPD\) Server table?](#)
- [Can SAS Scalable Performance Data \(SPD\) Server create indexes in parallel?](#)
- [Does SAS Scalable Performance Data \(SPD\) Server append indexes in parallel?](#)
- [What are ACLs and how do I use them to control access to data tables?](#)
- [How do I get a list of all the new SAS macro variables introduced for SAS Scalable Performance Data \(SPD\) Server?](#)
- [What about unique indexes? Can I do something to speed appends?](#)
- [What about disk compression for SAS Scalable Performance Data \(SPD\) Server tables?](#)
- [What about estimates for disk space consumption when using SAS Scalable Performance Data \(SPD\) Server?](#)
- [What should I set WORKPATH= to?](#)
- [How do I, as a LIBNAME domain owner, allow others to create tables in my domain?](#)
- [How does the system administrator list the access control lists for "user 1"?](#)
- [How do I change existing PROC SQL code that works with SAS to query SAS Scalable Performance Data \(SPD\) Server tables?](#)
- [Can I use pass-through async to create multiple indexes on a single existing table?](#)
- [Can I use pass-through async to create multiple indexes on existing tables?](#)
- [What size increases can I expect for tables that are stored in domains with BACKUP=YES?](#)

[Q. Does SAS Scalable Performance Data \(SPD\) Server support files greater than 2 Gigabytes?](#)

A. Yes. It does so by breaking up the files into partitions which are smaller than 2 Gigabytes. This is done automatically by the SAS Scalable Performance Data (SPD) Server host and requires no special syntax.

[Q. Can I create file systems greater than 2 Gigabytes?](#)

A. Yes, you can if you use a volume manager which will let you create file systems greater than 2 Gigabytes. This is our recommendation.

[Q. How do SAS Scalable Performance Data \(SPD\) Server client and server processes communicate?](#)

A. A SAS Scalable Performance Data (SPD) Server client communicates with three SAS Scalable Performance Data (SPD) Server processes.

When a client issues a LIBNAME assignment to the SAS Scalable Performance Data (SPD) Server host, the client communicates with the SAS Scalable Performance Data (SPD) Server Name Server process using the HOST= and SERV= options that were specified in the LIBNAME connection. The HOST= option specifies the host system where the SAS Scalable Performance Data (SPD) Server Name Server is running, and the SERV= option is the well-known port number of the SAS Scalable Performance Data (SPD) Server Name Server that was specified when was started. The SAS Scalable Performance Data (SPD) Server Name Server ensures that the domain of the LIBNAME assignment is valid, then returns the HOST= and SERV= option settings to the client. This ends the interaction of the client with the SAS Scalable Performance Data (SPD) Server Name Server for that LIBNAME assignment. The client will communicate with the SPDSSERV process to complete the LIBNAME assignment.

The SPDSSERV process authenticates the USER and PASSWORD portion of the LIBNAME assignment, and validates whether the USER has access to the domain. If the LIBNAME is successfully authenticated, the SPDSSERV process forks and executes a user proxy, the SPDSBASE process, which continues to service all other client requests for that LIBNAME connection. Subsequent LIBNAME assignments from the same client that are resolved to the same SAS Scalable Performance Data (SPD)

Server user and SPDSERV context are passed directly to SPDSBASE for processing without any further SPDSERV interaction. (No further interaction is required because the authentication is inherited by subsequent LIBNAME assignments.)

LIBNAME assignments from the same client for a different SAS Scalable Performance Data (SPD) Server user or LIBNAME assignments to a domain that is serviced by a different SPDSERV will result in a new SPDSBASE process to service that LIBNAME assignment.

Connections that use the record-level locking option LOCKING=YES to connect to a server in any domain are handled differently. All LIBNAME assignments share the same SPDSBASE record level locking process. When the LOCKING=YES option is in force, instead of forking and executing a new user proxy, the SPDSERV process initiates communication with the shared LOCKING=YES SPDSBASE process and the client.

Q. [How do I know which ports must be surfaced through an Internet firewall?](#)

A. There are two ports that the SAS Scalable Performance Data (SPD) Server Name Server uses that you can specify using command line options. The *listenport* option defines the port that must be used by clients (such as SAS) in LIBNAME and SQL CONNECT statements. The *listenport* option also can define the port that an ODBC data source requires to communicate with the SAS Scalable Performance Data (SPD) Server Name Server. The *operport* option defines a second port that is used for various command communications from SPDS utilities. Either of these ports may be specified using well-known port definitions in the operating system's services file, instead of specifying them on the command line. In UNIX systems this is typically the */etc/services* file. In the services file, the *spdsname* specification corresponds to *listenport*; and the *spdsoper* setting corresponds to the *operport* option. Both of these ports should be surfaced through the firewall.

The SPDSERV process uses two types of ports. The first type of port is a port that SPDSERV uses for local machine communications, internal to SAS Scalable Performance Data (SPD) Server. The second type of ports are ports that must be accessed by SAS Scalable Performance Data (SPD) Server clients.

Ports in the first category are not discussed here, because they do not need to be visible beyond the local machine, and therefore do not need firewall connectivity. There are two ports in the second category. The first port in the second category is the port that is defined by the SPDSERV *listenport* command line option that performs LIBNAME authentication of the SPDS user and password, and validates access to the SAS Scalable Performance Data (SPD) Server domain. The second port in the second category is the port that is used for various communications from SAS Scalable Performance Data (SPD) Server utilities, and is defined by the SPDSERV -*operport* command line option.

The SPDSERV -*listenport* and -*operport* specifications are registered in the SAS Scalable Performance Data (SPD) Server Name Server by the SPDSERV process when it starts up. Both specifications are returned to the SAS Scalable Performance Data (SPD) Server client from the SAS Scalable Performance Data (SPD) Server Name Server when it maps the LIBNAME domain to a SPDSERV. If you do not specify a *listenport* or *operport* in the SPDSERV command line, any port that is available will be used. Both of these ports should be specified in the SPDSERV command line and surfaced through the firewall.

Ports that the SPDSBASE process uses also fall into the two same categories. The first type of ports is used for local machine communications that are internal to SAS Scalable Performance Data (SPD) Server. The second type of ports are ports that must be accessed by SAS Scalable Performance Data (SPD) Server Clients. Like the SPDSERV process, the SPDSBASE process only cares about the ports that outside clients need to access through an Internet firewall.

The way that the SPDSBASE processes use ports is complex and requires a range of port numbers that are declared using the SAS Scalable Performance Data (SPD) Server MINPORTNO=MAXPORTNO= server parameter specifications. The MINPORTNO= and MAXPORTNO= parameters must both be specified in order to define the range of port numbers that are available to communicate with SAS Scalable Performance Data (SPD) Server clients, and therefore require access from outside of the firewall. If the SAS Scalable Performance Data (SPD) Server parameters for MINPORTNO= and MAXPORTNO= are not specified, the SPDSBASE processes will use any port that is available to communicate with the SAS Scalable Performance Data (SPD) Server client.

How many port numbers need to be set aside for SPDSBASE proxy processes? Each SPDSBASE process produces its own operator port that can be accessed using command-line specifications issued by a SAS Scalable Performance Data (SPD) Server client. In addition, each SAS Scalable Performance Data (SPD) Server table that is opened creates its own port. Each table's port becomes a dedicated data transfer connection that is used to stream data transfers to and from the SAS Scalable Performance Data (SPD) Server client. SAS Scalable Performance Data (SPD) Server table ports are normally dynamically assigned, unless the MINPORTNO= and MAXPORTNO= parameters have been specified. If the MINPORTNO= and MAXPORTNO= parameters have been specified, SAS Scalable Performance Data (SPD) Server table ports are assigned from within the specified port range.

Therefore, it follows that the range of ports that is specified for the MINPORTNO= and MAXPORTNO= parameters must take into consideration the peak number of concurrent LIBNAME connections that will be made to the server, as well as the I/O streams that are channelled between the SPDSBASE processes and the SAS Scalable Performance Data (SPD) Server clients.

The following ports **must** be surfaced for access beyond the firewall:

- Two SAS Scalable Performance Data (SPD) Server Name Server ports, -*listenport* and *operport*, as well as any other ports that are identified in SPDSNAME and SPDSOPER services.
- Two SPDSERV ports: -*listenport* and *operport*, as well as any other ports that are identified in SPDSERV_SAS and SPDSERV_OPER services.
- Any other ports that are defined in the MINPORTNO= and MAXPORTNO= range that is specified in the **spdserv.parm** file.

Q. [How does SAS Scalable Performance Data \(SPD\) Server interact with multi-homed hosts?](#)

A. A multi-homed host is a machine that has two or more IP addresses. For SAS Scalable Performance Data (SPD) Server to work properly on host machines with more than one IP address, you must define which IP address you want to associate with the socket bind calls that listen for the SAS Scalable Performance Data (SPD) Server Name Server and the SPDSERV processes. You use the SPDSBINDADDR environment variable to define the preferred IP address. You set the SPDSBINDADDR

environment variable in the **rc.spds** script that you use to initiate the SAS Scalable Performance Data (SPD) Server Name Server and SPDSSERV processes on the SAS Scalable Performance Data (SPD) Server host machine.

[Q. Can I use standard UNIX backup procedures?](#)

A. Yes. SAS Scalable Performance Data (SPD) Server files are standard files. If all the components of a table are in the same directory, then you can use the standard backup utility. This is our recommendation. SAS Scalable Performance Data (SPD) Server includes an incremental backup utility. The Help section on SAS Scalable Performance Data (SPD) Server Backup and Restore Utilities in the SAS Scalable Performance Data (SPD) Server Administrator's Guide contains additional information.

[Q. What do I need to know about SAS Scalable Performance Data \(SPD\) Server installation? How long will it take?](#)

A. The SAS Scalable Performance Data (SPD) Server install is quick and easy to do. The hardcopy installation instructions and shell scripts included on the install media will guide you through the installation process. Installation and verification take less than an hour. You may need additional time if you have a few SAS 8.2 client platforms to update.

On UNIX, the installation can be performed on a non-privileged UNIX account, although to implement all recommendations, UNIX root privilege is required. Consult the Installation section of the SAS Scalable Performance Data (SPD) Server Administrator's Guide for the SAS Scalable Performance Data (SPD) Server UNIX Installation Guide or the SAS Scalable Performance Data (SPD) Server Windows Installation Guide for additional information.

[Q. Is it necessary to run UNIX SAS Scalable Performance Data \(SPD\) Server as root?](#)

A. No. SAS recommends that you use a UNIX user ID *other than root* to run your production SAS Scalable Performance Data (SPD) Server environment. While there are no known security or integrity problems with the current SAS Scalable Performance Data (SPD) Server release, root access is not required to run the SAS Scalable Performance Data (SPD) Server environment when you properly configure the UNIX directory ownership and permissions on your LIBNAME domains. There is no real benefit from running the SAS Scalable Performance Data (SPD) Server package as root, other than possibly convenience. You should carefully consider whether any convenience you might obtain justifies the potential risk from running as root.

[Q. What is the Name Server and why do I need it?](#)

A. All access to SAS Scalable Performance Data (SPD) Server is controlled and managed by the Name Server. All clients first connect to the Name Server which acts as a gateway to named SAS Scalable Performance Data (SPD) Server domains. The Name Server maintains a dynamically updated list of valid SAS Scalable Performance Data (SPD) Server hosts and LIBNAME domains. When a user client needs a domain connection, the Name Server recurses the requested LIBNAME domain into a physical address and creates a proxy connection to the appropriate SAS Scalable Performance Data (SPD) Server host. Using the Name Server, the user does not have to know up-to-date physical addresses of SAS Scalable Performance Data (SPD) Server hosts. The only server that the client has to know about is the Name Server, which handles the details of connecting client users to appropriate domains.

[Q. Does every SAS Scalable Performance Data \(SPD\) Server client need a UNIX ID or Windows Networking ID?](#)

A. No. SAS Scalable Performance Data (SPD) Server does not use UNIX or Windows networking IDs for log in security. Each SAS Scalable Performance Data (SPD) Server client must have a valid SAS Scalable Performance Data (SPD) Server ID to log in to the server. Access to the server is controlled via this ID. Access to individual data is controlled by ACLs (Access Control Lists) which are created by the owner of the data.

[Q. Can a SAS Scalable Performance Data \(SPD\) Server host, Name Server and a SAS Scalable Performance Data \(SPD\) Server client all run on the same machine?](#)

A. Yes, they can. In fact, this even boosts performance because the client engine uses direct access where possible instead of issuing requests to the server. For example, the client will do direct reads from disk. WHERE-clause evaluation and index retrieval are faster, too.

[Q. Can I have multiple SAS Scalable Performance Data \(SPD\) Server hosts on the same machine?](#)

A. Yes. They can either be all connected to the same Name Server or different Name Servers. Within each Name Server, all SAS Scalable Performance Data (SPD) Server LIBNAME domains must be unique.

[Q. How do I create LIBNAME domains?](#)

-
A. LIBNAME domains are defined in a LIBNAME startup file. The required SAS Scalable Performance Data (SPD) Server command line option, -libnamefile, specifies the LIBNAME startup file. Each entry in this file has the form

```
LIBNAME=ldname pathname= . . . ;
```

where

LIBNAME = *ldname*, the libname domain name that clients reference

PATHNAME = *pathname*, the full UNIX or Windows path where data tables reside.

[Q. How do I specify a LIBNAME domain in SAS?](#)

A. LIBNAME domains are defined by using a SAS LIBNAME statement. A sample syntax is

```
libname sample sasspds 'ldname'  
server=spdshost.spdsname  
user='johndoe'  
prompt=yes ;
```

where

SAMPLE = name of LIBREF

SASSPDS = name of the SAS Scalable Performance Data (SPD) Server engine

LDNAME = LIBNAME Domain

SPDSHOST = the IP name of the node which is running the Name Server

SPDSNAME = port number which is being using by the Name Server

JOHNDOE = SAS Scalable Performance Data (SPD) Server login ID

PROMPT = prompt for password

[Q. Is there anything else I have to change to run my existing SAS applications?](#)

-
A. Typically, no. Once the LIBREFs have been assigned, your existing SAS application will run unchanged. However, there are some options that can be added to SAS programs that give additional performance boosts.

[Q. How can I get existing data loaded into a SAS Scalable Performance Data \(SPD\) Server table?](#)

-
A. There are several ways to accomplish this. Here are the three most common:

1. Use PROC COPY:

```
PROC COPY  
in=old  
out=spds  
memtype=data ;  
run ;
```

This will copy the data and build any existing indexes automatically.

2. Use the DATA Step and SET statement.

```
DATA spds.a ;  
set old.a ;  
run ;
```

This will copy the data. You have to specify the indexes to build.

```
DATA spds.a(index=(z));
  set old.a ;
run ;
```

This will copy the data and create an index on variable Z.

3. Use the Microsoft Windows ODBC driver.

Also, there is an excellent section in the SAS Scalable Performance Data (SPD) Server Help called [Migrating Tables Between SAS and SAS Scalable Performance Data \(SPD\) Server](#) which examines table conversions.

[Q. Can SAS Scalable Performance Data \(SPD\) Server create indexes in parallel?](#)

A. Yes, SAS Scalable Performance Data (SPD) Server can create multiple indexes at the same time. It does this by launching one thread per index and driving them all at the same time. You can accomplish this with

```
PROC DATASETS lib=spds ;
  modify a(asyncindex=yes) ;
  index create x ;
  index create y ;
  index create comp=(x y) ;
quit;
```

In the above example, X, Y and COMP will be created in parallel. Notice the ASYNCINDEX=YES data set option on the MODIFY statement.

```
%LET spdsiasy=YES ;
PROC DATASETS lib=spds ;
  modify a ;
  index create x ;
  index create y ;
  modify a ;
  index create
  comp=(x y)
  comp2=(y x) ;
quit ;
```

In the above example, X and Y will be created in parallel; COMP and COMP2 will be created in a second parallel index create as soon as the first pair completes. Notice the use of the SPDSIASY macro variable to specify parallel index creation. In this example, a table scan is required for each batch of indexes identified for creation in parallel: one table scan for the X and Y indexes and a second table scan for the COMP and COMP2 indexes.

How many indexes should you create in parallel? It will depend on how many CPUs are in the SMP configuration, available disk space for index key sorting, and other tasks. Some results show that on an 8-way UltraSparc, you can create 4 indexes in almost the same time it takes to create 1. You can group index creates to minimize table scans or auxiliary disk space consumption, but generally there is an inverse relationship between the two: minimizing table scans requires more auxiliary disk space and vice versa. The Help documentation contains more information on [parallel indexing](#).

[Q. Does SAS Scalable Performance Data \(SPD\) Server append indexes in parallel?](#)

A. Yes, SAS Scalable Performance Data (SPD) Server appends indexes in parallel by default.

[Q. What are ACLs and how do I use them to control access to data tables?](#)

A. ACLs are Access Control Lists which define who can access a data table and what kind of access they are granted. Currently, there are four levels of access defined: Access List Entry, Owner Access, Group Access, and Universal Access. Every SAS Scalable Performance Data (SPD) Server user has access to at least one group. During log in, a SAS Scalable Performance Data (SPD) Server user must specify a particular ACL group if the SAS Scalable Performance Data (SPD) Server password file has the user entered as a member of more than one group. Every data table has an ACL owner and the owner's ACL group attached to it. The precedence of the access levels is the following:

- Access List Entry
- Owner Access
- Group Access
- Universal Access

Types of access are READ, WRITE, ALTER and CONTROL. To create access lists you must have CONTROL access. The owner by default has control access. For more information, refer to the Help section in the SAS Scalable Performance Data (SPD) Server 4.4 Administrator's Guide on The ACL Command Set.

[Q. How do I get a list of the SAS macro variables introduced for SAS Scalable Performance Data \(SPD\) Server?](#)

A. In a SAS session, get into PROC SPDO and issue the SPDSMAC command. For example:

```
LIBNAME foo sasspds ... ;  
PROC SPDO lib=foo ;  
SPDSMAC ;
```

For more information, see the list of macro variables in the Help section on [SAS Scalable Performance Data \(SPD\) Server Macro Variables](#).

[Q. What about unique indexes? Can I do something to speed appends?](#)

A. You can use the [SPDSAUNQ=](#)YES server option to speed up appends to unique indexes.

[Q. What about disk compression for SAS Scalable Performance Data \(SPD\) Server tables?](#)

A. You can request compression for a SAS Scalable Performance Data (SPD) Server table by using the [COMPRESS=](#) data set option. You can also set a macro variable named [SPSDCMP](#) to the same value that you would set in the COMPRESS= option. This causes compression on all data sets you generate without explicitly specifying [COMPRESS=](#) on each DATA step. SAS Scalable Performance Data (SPD) Server compresses your table set by "blocks" and the way you control this amount is through the [IOBLOCKSIZE=](#) table option. Once you create a compressed table, the compression block size (that is, the number obs/block) cannot be changed. You must PROC COPY the data set to a new data set with a different [IOBLOCKSIZE=](#) on the output data set.

In any case, you select the default SAS Scalable Performance Data (SPD) Server compression by asserting COMPRESS=YES or using %let [SPSDCMP=](#)YES. The default compression algorithm is a run-length compression.

[Q. What about estimates for disk space consumption when using SAS Scalable Performance Data \(SPD\) Server?](#)

A. The answer to this question depends on what kind of component file within the SAS Scalable Performance Data (SPD) Server data you need to estimate. Recall that there are three classes of component files that make up a SAS Scalable Performance Data (SPD) Server table: metadata, data, and indexes. You always get the first two for every table. You get an index component file for each index you create on the table.

- [Metadata space consumption](#)
- [Data space consumption](#)
- [Hybrid index space consumption](#)
- [Transient space for PROC SORT/BY processing](#)

[Metadata space consumption](#)

The approximate estimate here is:

$$\text{SpaceBytes} = 12\text{Kb} + (\#\text{columns} * 120) + (5\text{Kb} * \#\text{indexes})$$

This estimate increases if you delete observations from the table or use compression on the table. In general, the size of this component file should not exceed approximately 400K.

[Data space consumption](#)

The estimate here is for uncompressed tables:

$$\text{SpaceBytes} = \#\text{rows} * \text{RowLength}$$

Your space consumption for compressed tables will obviously vary with the compression factor for your table as a whole.

[Hybrid index space consumption](#)

The hybrid index uses two data files. The .hbx file contains the global portion of the hybrid index. You can estimate space consumption roughly for the .hbx component of a hybrid index as follows:

If the index is NOT unique:

$$\text{number_of_discrete_values_in_the_index} * (22.5 + (\text{length_of_columns_composing_the_index}))$$

If the index IS unique:

$$\text{number_of_descrete_value_in_the_index} * (6 + (\text{length_of_columns_composing_the_index}))$$

The .idx file contains the per-value segment lists and bitmaps portion of the hybrid index. Estimating disk space consumption for this file is much more difficult than the .hbx file. This is because the .idx file size depends on the distribution of the key values across the rows of the table. The size also depends on the number of updates/appends performed on the index. The .idx files of an indexed table initially created with "n" rows consumes considerably less space than the .idx files of an identical table created and with several append or updates performed afterwards. The wasted space in the latter example can be reclaimed by reorganizing the index.

With the above in mind, a worst case estimate for space consumption of the .idx component of a hybrid index is:

$$8192 + (\text{number_of_descrete_values_in_more_than_one_obs} * (24 + (\text{avg_number_of_segments_per_value} * (16 + (\text{seg_size} / 8))))))$$

This estimate does not take into consideration the compression factor for the bitmaps, which could be substantial. The fewer occurrences of a value in a given segment, the more the bitmap for that segment can be compressed. The uncompressed bitmap size is the (seg_size/8) component of the algorithm.

To estimate the disk usage for non-unique hybrid index on a column with a length of 8, where the column contains 1024 discrete values, each value is in an average of 4 segments with a segment size of 8192 rows would be:

$$\text{.hyb_size} = 1024 * (22.5 + 8) = 31323 \text{ bytes}$$

$$\text{.idx_size} = 8192 + (10000 * (24 + (4 * (16 + (8192/8)))))) = 4343808 \text{ bytes}$$

To estimate the disk usage of a unique hybrid index on a column with a length of 8 that contains 100000 values would be:

$$\text{.hyb_size} = 100000 * (6 + 8) = 1400000 \text{ bytes}$$

$$\text{.idx_size} = 8192 + (0 * (...)) = 8192 \text{ bytes}$$

Note: The size of the .idx file for a unique index will always be 8192 bytes because the unique index contains no values that are in more than one observation.

There is a hidden workspace requirement when creating indexes or when appending indexes in SAS Scalable Performance Data (SPD) Server. This need arises from the fact that SAS Scalable Performance Data (SPD) Server sorts the rows of the table by the key value before adding the key values to the hybrid index. This greatly improves the index create/append performance but comes with a price requiring temporary disk space to hold the sorted keys while the index create/append is in progress. This work space is controlled for SAS Scalable Performance Data (SPD) Server by the WORKPATH= parameter in the SAS Scalable Performance Data (SPD) Server host parameter file.

You can estimate workspace requirements for index creation as follows for a given index "x":

$$\text{SpaceBytes} \sim \text{\#rows} * \text{SortLength}(x)$$

where

$$\text{\#rows} = \text{Number of rows in the table if creating; number of rows in the append if appending.}$$

if KeyLength(x) >= 20 bytes

$$\text{SortLength}(x) = (4 + \text{KeyLength}(x))$$

if KeyLength(x) < 20 bytes

$$\text{SortLength}(x) = 4 + (4 * \text{floor}((\text{KeyLength}(x) + 3) / 4))$$

For example, consider the following SAS code:

```
DATA foo.test ;
  length xc $15 ;
  do x=1 to 1000 ;
    xc = left(x) ;
    output ;
  end ;
run ;
```

```

PROC DATASETS lib=foo ;
  modify test ;
  index create x xc xxc=(x xc) ;
quit ;

```

For index X, space would be:

$$\begin{aligned}
\text{SpaceBytes} &= 1000 * (4 + (4 * \text{floor}((8 + 3) / 4))) \\
&= 1000 * (4 + (4 * \text{floor}(11 / 4))) \\
&= 1000 * (4 + 4 * 2) \\
&= 12,000
\end{aligned}$$

For index XC, space would be:

$$\begin{aligned}
\text{SpaceBytes} &= 1000 * (4 + (4 * \text{floor}(15 + 3) / 4))) \\
&= 1000 * (4 + (4 * \text{floor}(18 / 4))) \\
&= 1000 * (4 + 4 * 4) \\
&= 20,000
\end{aligned}$$

For index XXC, space would be:

$$\begin{aligned}
\text{SpaceBytes} &= 1000 * (4 + 23) \\
&= 1000 * 27 \\
&= 27,000
\end{aligned}$$

There is one other factor that plays into workspace computation: Are you creating the indexes in parallel or serially? If you create the indexes in parallel by using the [ASYNCINDEX=YES](#) data set option or by asserting the [SPDSIASY](#) macro variable, you will need to sum the space requirements for each index that you create in the same create phase. Referring back to the example in the parallel index Q/A, the indexes X and Y constitute a create phase, as do COMP and COMP2. You would need to sum the space requirement for X and Y, and for COMP and COMP2, and take the maximum of these two numbers to get the workspace needed to complete the PROC DATASETS indexes successfully.

The same applies to PROC APPEND runs when appending to the table with indexes. In this case all of the indexes are appended in parallel so you would need to sum the workspace requirement across all indexes.

Transient space for PROC SORT / BY processing

Workspace is required for SAS Scalable Performance Data (SPD) Server sorting just as it is required for SAS Scalable Performance Data (SPD) Server sorted index creation. There are two modes of sorting in SAS Scalable Performance Data (SPD) Server: tag and non-tag sorting. In either case you sort based on the columns selected in the BY clause. The difference is in the auxiliary data that is carried along by the sort in addition to the key constructed from the BY columns. The default for SAS Scalable Performance Data (SPD) Server is to use the non-tag sort.

In the case of non-tag sorting, SAS Scalable Performance Data (SPD) Server carries along the entire row contents (that is, all columns) as the auxiliary data for the key. In the mode of tag sorting, SAS Scalable Performance Data (SPD) Server only carries along the row ID that points back to the original table row as the auxiliary data. You control the amount of a sort problem that fits in memory at one time by the SAS Scalable Performance Data (SPD) Server parameter SORTSIZE. Obviously, for a given sort size the number of sort records that will fit will be a function of the sort mode($\#records = \text{SORTSIZE} / (\text{SortKeyLength} + \text{AuxillaryLength})$). When the sort problem does not fit in one SORTSIZE bin, the bins are written to workspace on disk and then merged back to make the final sorted run.

Estimating the disk space required for SAS Scalable Performance Data (SPD) Server sorting depends on the mode.

For non-tag sorting the estimate is

$$\text{SpaceBytes} = \#rows * (\text{SortKeyLength} + 4 + \text{RowLength})$$

For tag sorting the estimate is

$$\text{SpaceBytes} = \#rows * (\text{SortKeyLength} + 8)$$

So there is a very obvious question here: Since non-tag sort requires so much more space than a tag sort, why would you ever choose a non-tag sort, much less make it the default? The answer lies in the post-processing phase required for the tag sort. When the tag sort completes all you have is the sorted list of row IDs. You must probe the table using the row IDs to return the rows in the desired order. This generally means a highly randomized I/O access pattern to the original table which can add significantly to the time to complete the BY clause. There is definitely a trade-off between tag and non-tag sorting. The critical factors are the row length, the total number of rows to process, and the clustering of consecutive row IDs in the final ordering.

[Q: What should I set WORKPATH= to?](#)

A. The default server parameter file (that is, the spdserv.parm file) sets the WORKPATH= to "/var/tmp". Generally this will be inadequate for even moderate SAS Scalable Performance Data (SPD) Server usage. More than likely "/var/tmp" or "/tmp" at your site will have a small amount of space allocated and will be used by other system and application programs that create temporary files. In addition, these file systems will frequently be configured as a memory mapped file system by the system administrator (that is, mounted as a tmpfs file system). Our experience has been that neither "/var/tmp" nor "/tmp" is a suitable choice for WORKPATH given the space and performance limitations of a typical system configuration.

We strongly recommend that you configure WORKPATH to use a volume manager file system that is striped across multiple disks if possible to provide the optimum performance and to allow adequate temporary workspace for the collection of SAS Scalable Performance Data (SPD) Server proxy processes that will be running concurrently on your server hardware.

[Q: How do I, as a LIBNAME domain owner, allow others to create tables in my domain?](#)

A: For example, Tom is a LIBNAME domain owner, and he wants to give Fred access to create tables in Tom's domain. Tom needs to do the following:

Description of the LIBNAME statement.

- **dmowner** is the LIBREF for the location of the SAS Scalable Performance Data (SPD) Server data.
- **tomdom** is the previously established SAS Scalable Performance Data (SPD) Server domain.
- **host=** specifies the name of the computer where SAS Scalable Performance Data (SPD) Server resides.
- **serv=** is followed by the port number of the SAS Scalable Performance Data (SPD) Server's Name Server.
- **passwd=** is followed by the required password for **tom**.

SAS code, by line	Remarks
<pre>LIBNAME dmowner sasspds "tomdom" host="samson" serv="5555" user="tom" passwd="tompw" ;</pre>	
<pre>PROC SPDO lib=dmowner ;</pre>	PROC SPDO opens the command set that allows the user tom to change ACLs in the tomdom domain using the LIBREF dmowner .
<pre>set acluser tom ;</pre>	SET ACLUSER command allows ACLs under user ID tom to be modified.
<pre>add acl/libname ;</pre>	Command to add the ACL for a LIBNAME domain. LIBNAME is the syntax used to indicate the LIBNAME domain assigned, which is tomdom to the LIBREF that PROC SPDO is started with, which is dmowner .
<pre>modify acl/LIBNAME fred=(Y,Y,,) ;</pre>	Modifies the ACL in the LIBNAME domain ACL to give user ID fred read and write access to the tomdom domain.
<pre>quit ;</pre>	

Fred can now connect to the TOMDOM domain and create tables.

[Q: How does the system administrator list the access control lists for "user 1"?](#)

To see the ACL privileges for a domain, the system administrator lists them for each user.

A: For this to work, your SAS Scalable Performance Data (SPD) Server user ID must be previously set up to have the SPECIAL (level 7) privilege, to use the ACLSPECIAL=YES option on a LIBNAME statement.

Command from command prompt >	Remark
<pre>LIBNAME test saspbs 'temp' server=servname.7880 prompt=yes ;</pre>	Issue LIBNAME statement for test domain, specify server and port number, ask system for a password prompt.
<pre>user="username" aclspecial=YES ;</pre>	aclspecial=YES now gives "username" access to special ACL commands, such as setting a new user ID.
<pre>PROC SPDO lib=test ;</pre>	Connects to the temp LIBNAME domain using the LIBREF test .
<pre>set acluser user1 ;</pre>	Sets the SAS Scalable Performance Data (SPD) Server user scope to user1 .
<pre>list acl _all_ ;</pre>	Lists all ACLs owned by user1 .

The resulting output, described in the table below, lists all of the tables in "test".

Resulting output from list acl _all_ ; command	Remarks
The SAS System 10:58 Tuesday, November 17, 2003	System message
ACL Info for A.DATA	This ACL will affect table A if table A exists and user1 is the owner or has ACL control of the table A.
Owner = USER1	USER1 created and owns the A.DATA ACL.
Group = TECH	This ACL was created while user1 was connected with an ACL group of TECH. All group permissions will affect the permissions of the members of the TECH ACL group.
Default Access (R,W,A,C) = (Y,N,N,N)	R=Read; W=Write; A=Alter (rename, delete, or replace tables) C=Control (define and update ACLs for a table) Y=Yes; N=No; Universal privileges are limited to read on table A.DATA.
Group Access (R,W,A,C) = (N,N,N,N)	Users in the ACL group TECH have no privileges on table A.DATA.
The SAS System 10:58 Tuesday, November 17, 2003	
ACL Info for NTE*.DATA	NTE*.DATA refers to a set of tables, which begin with NTE. ACLs of this kind are created using the generic option. If you create a specific ACL for a table that starts with NTE, the specific ACL will override the generic ACL.
Owner = user1	
Group = TECH	

Default Access (R,W,A,C) = (N,N,N,N)	
Group Access (R,W,A,C) = (Y,Y,N,N)	Users from the ACL group TECH have read and write access to tables with names that start with NTE.

Q: [How do I change existing PROC SQL code that works with SAS to query SAS Scalable Performance Data \(SPD\) Server tables?](#)

A: You do not have to change your PROC SQL code. The way to do this is to wrap your code inside a CONNECT statement, which points to the location of the SAS Scalable Performance Data (SPD) Server tables. This technique is referred to as pass-through. Normal operating system and ACL privileges apply to the user ID making the query during the CONNECT process. Your PROC SQL code should work with a few [exceptions](#).

Once you establish a working CONNECT statement which points to the location of your SAS Scalable Performance Data (SPD) Server tables, you can assign a LIBNAME to the SAS Scalable Performance Data (SPD) Server table path with a LIBREF command. This allows the simple name you assign to the SAS Scalable Performance Data (SPD) Server table to be used in the SQL query, which keeps your SQL query as short as possible.

Here are four progressive examples:

- [The first example](#) shows PROC SQL that works with SAS.
- [The second example](#) shows how you can access SAS Scalable Performance Data (SPD) Server tables without using the pass-through facility.
- [The third example](#) shows how you can access the SAS Scalable Performance Data (SPD) Server tables by changing your PROC SQL code.
- [The fourth example](#) shows how you can use the original PROC SQL code from the first example, wrapped with a CONNECT statement, so that it can query SAS Scalable Performance Data (SPD) Server tables.

Example 1: PROC SQL query, designed to work with a SAS data set, with a two-level SAS filename example.

Code	Remarks
<pre> /* Issue a LIBNAME statement which */ /* creates a LIBREF called "test" */ LIBNAME test '/path/for/your/data' ; /* Query using base LIBREF of test */ PROC SQL ; select sum(table1+table2) as pass, carrier from test.carriers where carrier in ('AA','JI') and bstate='TX' group by carrier ; quit ; </pre>	<p>This is an example of SQL code that works with SAS. The code contains a two-level SAS file name reference, which is typical for PROC SQL, but it will not work if we attempt to use it inside of a pass-through CONNECT statement.</p> <p>Each of the following examples will show variations of this code, modified to access SAS Scalable Performance Data (SPD) Server information. We'll also discuss the pros and cons of each method.</p>

Example 2: PROC SQL query, without using pass-through, pointing to a SAS Scalable Performance Data (SPD) Server table, with a two-level SAS file name.

Why would you want to do this? You might NOT want to do this, because without pass-through, all the processing is done on the CPU of the client machine. When processing large tables, this is impractical, if not impossible.

Code	Remarks:


```

/* Issue a SAS Scalable Performance Data (SPD) Server (mkt) library      */
/* LIBNAME statement                                                    */

LIBNAME mkt sasspds 'mkt'
  server=servername.4228
  user='anonymous' ;
PROC CONTENTS data=mkt.carriers ;
  run ;

/* query spds LIBREF (mkt)                                             */
/* (two-level SAS file name)                                          */

PROC SQL;
  select sum(table1+table2)
  as pass, carrier
  from mkt.carriers
  where carrier in
  ('AA','JI')
  and bstate='TX'
  group by carrier ;

quit ;

```

This example shows you how to make a query against SAS Scalable Performance Data (SPD) Server tables, using your original SQL code, without using SQL pass-through.

Example 3: PROC SQL query, using pass-through, pointing to a SAS Scalable Performance Data (SPD) Server table, with a LIBNAME example, with SQL code modified, to avoid using a two-level SAS file name.

Why would you want to do this? By modifying your SQL code slightly, you can use the pass-through facility to have SAS Scalable Performance Data (SPD) Server perform the work and send the results to the client.

Code

```

/* Query spds LIBREF (mkt) (pass-through one-level LIBREF )*/

PROC SQL;

connect to sasspds
  (dbq='mkt'
  serv='8770'
  user='anonymous'
  host='localhost') ;

select *
  from connection
  to sasspds

(select sum(table1+table2)
  as pass,
  carrier
  from carriers
  where carrier
  in('AA','JI')
  and bstate='TX'
  group by carrier) ;

quit ;

```

Example 4: PROC SQL query, using pass-through, pointing to a SAS Scalable Performance Data (SPD) Server table, executing a LIBREF statement on the server, so that existing code can be used "as is."

Why would you want to do this?

Without modifying your SQL code, you can use the pass-through facility so SAS Scalable Performance Data (SPD) Server performs the work and sends the results to the client.

Code

```
PROC SQL ;

connect to sasspds
  (dbq='mkt'
   serv='8770'
   user='anonymous'
   host='localhost');

/* Issue passthru LIBREF (mkt) for use */
/* in two-level queries                */

execute(LIBREF mkt)
  by sasspds;

/* Query the SAS Scalable Performance Data (SPD) Server LIBREF (mkt) */
/* that is a pass-through LIBREF      */

select *
  from connection
  to sasspds

(select sum(table1+table2)
 as pass,
 carrier from mkt.carriers
 where carrier in('AA','JI')
 and bstate='TX'
 group by carrier) ;

quit;
```

[Q: Can I use pass-through async to create multiple indexes on a single existing table?](#)

A: No. Multiple create indexes on the same existing table are not supported with async.

PROC DATASETS can be used to create indexes in parallel for a single existing table.

For example:

```
PROC DATASETS lib=foo ;
modify customer(
  asyncindex=yes
  bitindex=(state) ;
index create state ;
index create phoneno ;
index create custno ;
index create totsales ;
quit ;
```

[Q: Can I use pass-through async to create multiple indexes on existing tables?](#)

A: Yes. As long as you create only one index per table, the index creation can be run with async.

For example, to create an index State on table Customer, an index Totals on table Billing, and an index Orderno on table Orders asynchronously, you use the following code:

```
execute(begin async operation)
  by sasspds ;

execute(create index state on customer(state))
```

```
by sasspds ;

execute(create index totals on billing(totals))
by sasspds ;

execute(create index orderno on orders(orderno))
by sasspds ;

execute(end async operation)
by sasspds ;
```

Q: [What size increases can I expect for tables that are stored in domains with BACKUP=YES?](#)

A: Tables created in domains that have Backup=YES will have an additional 17 bytes per observation.

[Copyright © 2008 by SAS Institute Inc., Cary, NC, USA. All rights reserved.](#)

SAS Scalable Performance Data (SPD) Server Operator Interface Procedure (PROC SPDO)

- [Special SPDO Commands](#)
 - [SPDO Command Examples](#)
 - [LIBNAME Proxy Commands](#)
 - [LIBNAME Proxy Command Examples](#)
 - [Privileged OPER Commands](#)
 - [TRUNCATE Command and Example](#)
 - [Refreshing SPD Server Parameter and LIBNAME Files](#)
 - [REFRESH Command Examples](#)
 - [Commands to Nonexistent Users](#)
-

PROC SPDO

Note: For additional information on using PROC SPDO, see the chapter, "Controlling SAS Scalable Performance Data Server Resources with PROC SPDO and ACL Commands," in the *SAS Scalable Performance Data Server: Administrator's Guide*.

[Special SPDO Commands](#)

The following SPDO commands require that you have ACLSPECIAL= enabled for your SPDO LIBNAME connection.

To enable ACLSPECIAL=, you must first grant the SPD Server user ID ACLSPECIAL= access rights. Second, the user must request access for a specific connection. To request access, the user adds the ACLSPECIAL=YES option to the LIBNAME statement. The user can now submit the following SPDO commands:

```
SPDSCMD 'command'
```

This example runs the specified command in the context of the user ID for the spdserv process that is associated with the LIBNAME connection that the user used to invoke PROC SPDO. No restrictions are placed on commands that are executed in this manner. Therefore, you must carefully consider which SPD Server users need ACLSPECIAL access rights.

[SPDO Command Examples:](#)

1. List the WORKPATH directory:

```
spds cmd 'ls /spdswork/*.spds';  
spds cmd 'dir d:\spdswork\*.spds';
```

2. Clean up WORKPATH files:

```
spdscmd 'rm /spdswork/*.spds';  
spdscmd 'del d:\spdswork\*.spds';
```

LIBNAME Proxy Commands

To issue proxy commands, you must first select the SPD Server user proxy.

LIST USERS;

lists the proxy processes that are accessible to the PROC SPDO *lib=<libname>* statement that was dispatched from the SPD Server host. Accessible proxies are anonymous proxies and proxies that are owned by the LIBNAME owner. If the LIBNAME owner has ACLSPECIAL privileges, then all user proxies will be listed.

SET USER *userID* [*portnumber*];

allows you to use the port number to distinguish between two proxies that share the same user ID.

LIST USERS/LOCKING;

lists the user-locking proxy threads that are accessible by the PROC SPDO *lib=<libname>* statement that was dispatched from the SPD Server host and that were created with the LOCKING=YES LIBNAME option. Accessible proxies are anonymous proxies and proxies that are owned by the LIBNAME owner. If the LIBNAME owner has ACLSPECIAL privileges, then all user-locking proxies will be listed. For each user-locking proxy thread, SPD Server returns the SPD Server user ID, the client login, and the thread ID. You can select a user-locking proxy thread from the LIST USERS list by submitting a command in the following form:

SET USER/LOCKING [*userID threadID=#*];

Once a user-locking proxy is selected, you can get LIBNAME information by submitting the following commands:

```
SHOWLIBNAME libref | _ALL_;
```

```
SHOWLIBNAME libref / DATA= [ _ALL_ / dsname];
```

```
SHOWLIBNAME libref / DUMP= [ _ALL_ / dsname];
```

where *libref* is an explicit SPD Server LIBNAME name. Specify *_ALL_* to see every currently assigned LIBNAME for the proxy.

If the /DATA= option is used with *_ALL_*, information about all of the open tables in the proxy for the given LIBNAME is displayed. If the /DATA= option is used with a data set name *dsname*, detailed information about the specified data set table is displayed.

If the /DUMP= option is used with *_ALL_*, information about all of the accessible tables in the proxy for the given LIBNAME is displayed. If the /DATA= option is used with a data set name *dsname*, detailed information about the specified data set table is displayed.

LIBNAME Proxy Command Examples

1. List all of the users for the server 'sunburn.6100':

```
LIBNAME example sasspds
  host='sunburn'
  serv='6100'
  user='sassyl'
  passwd='abc123'
  aclspecial=YES;
```

```
PROC SPDO lib=example;
list users;
```

Users Currently Connected to SPD Server

UserName Pid Portno

```
-----
SASSYL    17704 58382
SASSYL    17614 58298
SASSYL    17613 58293
ANONYMOU  17611 58288
ANONYMOU  17610 58283
```

2. Set the user to ANONYMOU and specify process ID (Pid) 17610:

```
set user anonymou 17610;
```

NOTE: User ANONYMOU connected to proxy operator port with pid=17610.

3. Show every LIBNAME for user ANONYMOU for this proxy:

```
showlibname _all_;
LIBREF(FOO):Pathname assigned=/bigdisk/test/qabigl_dev/
LIBREF(FOO):ACL Owner=
LIBREF(FOO):ACL Defaults(R,W,A,C)=(Y,Y,Y,Y)
```

4. Show all of the open tables in LIBNAME FOO:

```
showlibname FOO/data=_all_;
```

NOTE: No data sets currently opened for LIBREF FOO.

5. Show all of the accessible tables in LIBNAME FOO:

```
showlibname FOO/dump=_all_;
```

```
LIBREF(FOO):Dataset name=BIGX  
LIBREF(FOO):ACL Owner=ANONYMOU  
LIBREF(FOO):ACL Defaults(R,W,A,C)=(N,N,N,N)  
LIBREF(FOO):Dataset name=X  
LIBREF(FOO):ACL Owner=ANONYMOU  
LIBREF(FOO):ACL Defaults(R,W,A,C)=(N,N,N,N)
```

6. The user ANONYMOU performs a WHERE clause on the table BIGX. Show all of the open tables in LIBNAME FOO:

```
showlibname FOO/data=_all_;
```

```
LIBREF(FOO):Dataset name=BIGX  
LIBREF(FOO):ACL Owner=ANONYMOU  
LIBREF(FOO):ACL Defaults(R,W,A,C)=(N,N,N,N)  
LIBREF(FOO):WHERE clause read thread active
```

7. User ANONYMOU performs a WHERE clause on the table BIGX and displays detailed information about the table BIGX:

```
showlibname FOO/data=bigx;
```

```
LIBREF(FOO):Dataset name=BIGX  
LIBREF(FOO):ACL Owner=ANONYMOU  
LIBREF(FOO):ACL Defaults(R,W,A,C)=(N,N,N,N)  
LIBREF(FOO):WHERE clause read thread active  
LIBREF(FOO):Type=  
LIBREF(FOO):Label=  
LIBREF(FOO):Number observations=5000000  
LIBREF(FOO):Observation length=41  
LIBREF(FOO):Wire blocksize=32718  
LIBREF(FOO):Wire block factor=798  
LIBREF(FOO):Data port number=58392  
LIBREF(FOO):Active data socket=33  
LIBREF(FOO):Metafile=/bigdisk/test/qabig1_dev/bigx.mdf.0.0.0.spds9  
LIBREF(FOO):Metafile size=31  
LIBREF(FOO):Datafile=  
  /spds02/test/qabig1_dev/bigx.dpf._bigdisk_test_qabig1_dev.0.1.spds9:  
  /spds03/test/qabig1_dev/bigx.dpf._bigdisk_test_qabig1_dev.1.1.spds9:  
  /spds04/test/qabig1_dev/bigx.dpf._bigdisk_test_qabig1_dev.2.1.spds9:  
  /spds01/test/qabig1_dev/bigx.dpf._bigdisk_test_qabig1_dev.3.1.spds9:  
  /spds02/test/qabig1_dev/bigx.dpf._bigdisk_test_qabig1_dev.4.1.spds9:  
  /spds03/test/qabig1_dev/bigx.dpf._bigdisk_test_qabig1_dev.5.1.spds9:  
  /spds04/test/qabig1_dev/bigx.dpf._bigdisk_test_qabig1_dev.6.1.spds9:  
  /spds01/test/qabig1_dev/bigx.dpf._bigdisk_test_qabig1_dev.7.1.spds9:  
  /spds02/test/qabig1_dev/bigx.dpf._bigdisk_test_qabig1_dev.8.1.spds9:
```

```
/spds03/test/qabig1_dev/bigx.dpf._bigdisk_test_qabig1_dev.9.1.spds9:
/spds04/test/qabig1_dev/bigx.dpf._bigdisk_test_qabig1_dev.10.1.spds9:
/spds01/test/qabig1_dev/bigx.dpf._bigdisk_test_qabig1_dev.11.1.spds9:
/spds02/test/qabig1_dev/bigx.dpf._bigdisk_test_qabig1_dev.12.1.spds9
LIBREF(FOO):Datafile size=200196
LIBREF(FOO):Number of Indexes=0
```

8. List all locking users for the server 'sunburn.6100':

```
list users/locking;
```

```
Users Currently Connected to the Record Level Proxy
```

```
SPDUserName Client Login Thread Id
```

```
-----
ANONYMOU     SASTEST      7
TEST         SASTEST      8
```

9. Set the user to ANONYMOU and specify thread ID 7:

```
set user/locking anonymou threadid 7;
```

```
NOTE: User ANONYMOU connected to record
level proxy operator port with thread=7.
```

10. Show every LIBNAME for locking user ANONYMOU:

```
showlibname _all_;
```

```
LIBREF(LOCKING):Pathname assigned=/bigdisk/test/qabig1/
LIBREF(LOCKING):ACL Owner=
LIBREF(LOCKING):ACL Defaults(R,W,A,C)=(Y,Y,Y,Y)
```

11. Show all of the open tables in LIBNAME LOCKING:

```
showlibname LOCKING/data=_all_;
```

```
NOTE: No data sets currently opened for LIBREF LOCKING.
```

Privileged OPER Commands

You must have ACLSPECIAL access rights (LIBNAME option ACLSPECIAL=YES) to run privileged OPER commands. With privileged OPER commands, you must first set yourself as the proxy operator by

submitting the following command:

SET MODE OPER;

The SET MODE OPER command sets you as the operator of the user proxy that you are currently set to. There can be only one operator for a user proxy at any time. If you submit the SET MODE OPER command when someone is already established as operator of the user proxy, you get the following message:

```
ERROR: Operator mode owned by another connection.  
Cannot grant this request.
```

After you have successfully set yourself as the operator, the following commands can be submitted:

OPER CANCEL [/DUMP];

The OPER CANCEL command cancels and exits the user proxy. If the /DUMP option is specified for a non-locking user proxy, the proxy exits with an abort() call, which produces a core file. If you are the operator of a locking user proxy, the /DUMP option is ignored. The OPER CANCEL command initiates a hard exit of the user proxy. Hard exits might leave tables opened for UPDATE access, which is an inconsistent and unusable state. In this case, you can submit the PROC DATASETS REPAIR command to restore the tables to a usable state.

OPER DISCONNECT;

The OPER DISCONNECT command drops the control socket from the user proxy to the client. This action causes the user proxy to terminate the next time it tries to communicate with the client. This termination initiates a hard exit of the user proxy. Hard exits might leave tables opened for UPDATE access, which is an inconsistent and unusable state. In this case, you can submit the PROC DATASETS REPAIR command to restore the tables to a usable state.

The OPER DISCONNECT command differs from the OPER CANCEL command. When the OPER DISCONNECT command is submitted, the user proxy continues until it detects that the control socket connection has been dropped. As a result, the OPER DISCONNECT command has the potential to complete. However, when the control socket disconnect is detected by the user proxy varies, with different results.

OPER INTERRUPT;

The OPER INTERRUPT command sets a soft interrupt flag in any open tables that belong to the user proxy. During certain long-running operations, such as large table sorts, table scans with a WHERE clause, or index creations, the user proxy periodically checks for an interrupt flag in all of the open tables that are involved in the operation. If an interrupt flag is detected, the user proxy terminates the operation and any previously opened tables are closed.

Unlike the OPER CANCEL command or the OPER DISCONNECT command, the OPER INTERRUPT command initiates a soft exit of the user proxy. The user receives a message in the SAS log that states that the job has been interrupted. If the job did not finish, then the results might be incomplete. However, the user LIBNAME will be intact, and the user proxy will still be viable.

Whether a job will be interrupted cannot be determined; it depends on the job that is currently running. To determine if a job can be interrupted, submit a SHOWLIBNAME *libref/*

DATA=_ALL_ command before you submit the OPER INTERRUPT command to see all of the open tables. You can also submit the SHOWLIBNAME *libref*/ DATA=_ALL_ command after you submit the OPER INTERRUPT command, to see if all of the open tables were closed. If the tables are still open after the OPER INTERRUPT command has been submitted, you can wait and check again later. If the tables need to be closed immediately, you can use the OPER CANCEL command to cancel the user proxy.

TRUNCATE Command and Example

The TRUNCATE command is a PROC SPDO command that allows you to delete all of the rows in a table without deleting the table structure or metadata.

```
%let host=kaboom;
%let port=5191;
%let domain=path2;

LIBNAME &domain sasspds "&domain"
  server=&host..&port
  user='anonymous'
  ip=YES;

/* create a table */

DATA &domain..staceys_table;

  do i = 1 to 100;
    output;
  end;
run;

/* verify the contents of the created table */

PROC CONTENTS data=&domain..staceys_table ;
run;

/* SPDO Truncate command deletes the table */
/* data but leaves the table structure in */
/* place so new data can be appended */

PROC SPDO lib=&domain;
SET acluser;
TRUNCATE staceys_table;

quit;

/* verify that no rows or data remain in */
```

```
/* the structure of staceys_table          */  
  
PROC CONTENTS data=&domain..staceys_table;  
run;
```

Refreshing SPD Server Parameter and LIBNAME Files

You can use PROC SPDO to dynamically refresh SPD Server parameter and LIBNAME files. If you make changes to your spdsserv.parm file or to your libnames.parm environment file for SPD Server, you can use the REFRESH command to avoid restarting the server. Submitting the REFRESH command refreshes the specified SPD Server file without restarting the server.

When you submit the REFRESH command, SPD Server refreshes the operating parameter file.

The syntax for the REFRESH command to refresh the libnames.parm file is:

REFRESH DOMAINS

The syntax for the REFRESH command to refresh the spdsserv.parm file is:

REFRESH PARMS

Each REFRESH operation completely refreshes and replaces the contents of the previous libnames.parm file or the spdsserv.parm file in the SPD Server environment.

REFRESH Command Examples

Add a new LIBNAME domain to your current libnames.parm file and use it without restarting the server:

```
LIBNAME spds44 sasspds 'spds44'  
    server=estore.5180  
    user='admin'  
    password='spds123'  
    aclspecial=YES  
    prompt=YES;  
  
PROC SPDO library=spds44;  
    SET acluser admin;  
    REFRESH PARMS;  
    REFRESH DOMAINS;  
quit;
```

Here is a more detailed example:

```
/* Domain reftest is a pre-existing domain. */
```

```

/* Add domain reftest2 to libnames.parm and */
/* specify owner=admin */

LIBNAME=tmp pathname=c:\temp;
LIBNAME=formats pathname=c:\data\formats;
LIBNAME=reftest pathname=c:\data\reftest
  owner=admin;
LIBNAME=reftest2 pathname=c:\data\reftest2
  owner=admin;

/* Run refresh job using admin with ACLSPECIAL */
/* The SPD Server user must have ACLSPECIAL */
/* privileges to refresh domains. */

LIBNAME reftest sasspds 'reftest'
  server=d8488.5180
  user='admin'
  password='spds123'
  aclspecial=YES;

PROC SPDO library=reftest;
  SET acluser admin;
  REFRESH DOMAINS;
quit;

/* Domains that have an owner= option such as */
/* reftest2 (owner=admin) must be reconnected */
/* to the domain again. */

LIBNAME reftest2 sasspds 'reftest2'
  server=d8488.5180
  user='admin'
  password='spds123';

```

Commands to Nonexistent Users

In SPD Server, you can submit operator commands to a user after selecting the user with the SET USER or SET USER/RECORD command. However, user sessions are finite. The user that you select with SET USER or SET USER/RECORD might be unavailable when the user ends the SAS session, or disconnects from a LIBNAME and the user proxy. If you submit an OPER command to a user that is no longer in session, or to a user that has ended a locking user proxy, you get the following message:

```
ERROR: Specified locking user no longer exists.
```

If the disconnected user used a non-locking user proxy, and you submit an OPER command, you get the following

message:

ERROR: Specified user <userID> with pid <Process-ID> no longer exists.

Either of these messages indicates that the user that was selected is no longer valid. In this case, you must use SET USER or SET USER/RECORD to select a different user.

Copyright © 2008 by SAS Institute Inc., Cary, NC, USA. All rights reserved.

SAS Scalable Performance Data (SPD) Server STARJOIN Facility

- [Overview: SPD Server STARJOIN Facility](#)
- [Star Schemas](#)
- [SPD Server STARJOIN Requirements](#)
- [Enabling STARJOIN Optimization](#)
- [Invoking the SPD Server STARJOIN Facility](#)
- [SPD Server STARJOIN Optimization](#)
 - [Classify Dimension Tables That Are Called by SQL as Phase I Tables or Phase II Tables](#)
 - [Phase I Probes Fact Table Indexes and Selects a STARJOIN Strategy](#)
 - [Phase II Performs Index Lookups and Joins Subsetted Fact Table Rows with Phase II Tables](#)
- [Indexing Strategies to Optimize STARJOIN Query Performance](#)
 - [Indexing to Optimize the IN-SET Join Strategy](#)
 - [Indexing to Optimize the COMPOSITE Join Strategy](#)
 - [Example: Indexing using IN-SET Join Strategy](#)
- [SPD Server STARJOIN RESET Statement Options](#)
 - [Example: STARJOIN RESET Statement Options](#)
- [SPD Server STARJOIN Examples](#)
 - [Example 1: Valid SQL STARJOIN Candidate](#)
 - [Example 2: Invalid SQL STARJOIN Candidate](#)
 - [Example 3: STARJOIN Candidate with Created or Calculated Columns](#)

SPD Server STARJOIN Facility Figures

- [Figure 1: Example Star Schema](#)
-

[Overview: SPD Server STARJOIN Facility](#)

SPD Server provides an enhanced SQL planner that includes the STARJOIN facility. The SPD Server STARJOIN facility validates, optimizes, and executes SQL queries on data that is configured in a star schema. Star schemas are composed of two or more normalized dimension tables that surround a centralized fact table. The centralized fact table contains data elements of interest that are derived from the dimension tables.

In data warehouses with large numbers of tables and millions or billions of rows of data, a properly constructed STARJOIN can minimize overhead data redundancy during query evaluation. If the SPD

Server STARJOIN facility is not enabled, or if SPD Server SQL does not detect a star schema, then the SQL will be processed using pairwise joins.

How does a STARJOIN differ from a pairwise join? In SPD Server, a properly configured STARJOIN requires only three steps to complete, regardless of the number of dimension tables. SPD Server pairwise joins require one step for each table to complete the join. If a star schema consists of 25 dimension tables and one fact table, the STARJOIN is accomplished in three steps; joining the tables in the star schema using pairwise joins requires 26 steps.

When data is configured in a valid SPD Server star schema, and the STARJOIN facility is not disabled, the SPD Server STARJOIN facility can produce quicker and more processor-efficient SQL query performance than SQL pairwise joins.

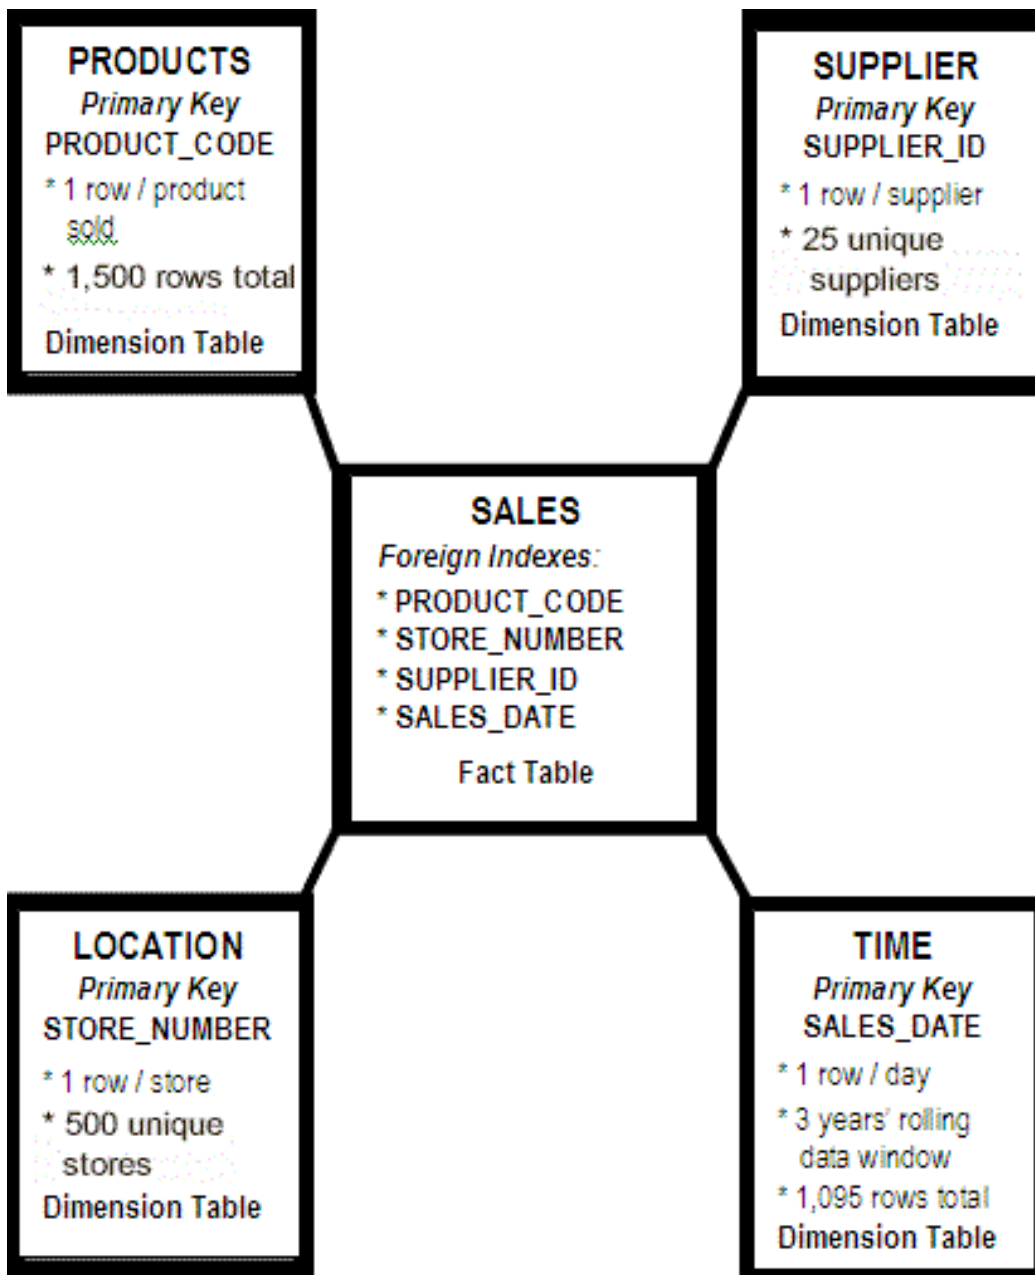
Star Schemas

To exploit the power of the SPD Server STARJOIN facility, the data must be configured as a star schema, and it must meet specific SPD Server SQL star schema requirements.

Star schemas are the simplest data warehouse schema, consisting of a central fact table that is surrounded by multiple normalized dimension tables. Fact tables contain the measures of interest. Dimension tables provide detailed information about the attributes within each dimension. The columns in fact tables are either foreign key columns that define the links between the fact table and individual dimension tables, or they are columns that calculate numeric values that are based on foreign key data.

Figure 1 is an example of a simple star schema. The dimension tables Products, Supplier, Location, and Time surround the fact table Sales.

Figure 1: Example Star Schema



The dimension tables, fact table, and keys in Figure 1 are used in the examples in this document.

Dimension Tables Information

Products is a table of products, with one row per unique product SKU. The row for each unique SKU contains information such as product name, height, width, depth, weight, pallet cube, and so on. The example Products table contains 1,500 rows.

Supplier is a table of the suppliers that supply the products. The row for each unique supplier contains information such as supplier name, address, state, contact representative, and so on. The example Supplier table contains 25 rows.

Location is a table of the stores selling the products. The row for each unique location

contains information such as store number, store name, store address, store manager, store sales volume, and so on. The Location table contains 500 rows.

Time is a sequential sales transaction table. Each row in the Time table represents one day out of a rolling three-year, 365-day-per-year calendar. The row for each day contains information such as date, day of week, month, quarter, year, and so on. The Time table contains 1,095 rows.

Fact Table Information

The fact table Sales is a table that combines information from the four dimension tables, Products, Supplier, Location, and Time. Its foreign keys are imported, one from each dimension table: PRODUCT_CODE from Products, STORE_NUMBER from Location, SUPPLIER_ID from Supplier, and SALES_DATE from Time. The fact table Sales might have other columns with facts or information that are not found in any dimension table. Examples of fact table columns that are not foreign keys from a dimension table are columns such as QTY_SOLD or NET_SALES. The fact table in this example could contain as many as $1,500 \times 25 \times 500 \times 1,095 = 20,531,250,000$ rows.

SPD Server STARJOIN Requirements

For SPD Server SQL to take advantage of the STARJOIN planner, the following conditions must be true:

- STARJOIN optimization must be enabled in SPD Server.
- The SPD Server star schema must use a single central fact table.
- All dimension tables in the SPD Server star schema must be connected to the fact table.
- SPD Server dimension tables can appear in only one join condition.
- SPD Server fact tables are equally joined to dimension tables.
- SPD Server SQL infers fact tables by topology (common equally joined predicates).
- Dimension tables that have no subsetting require a simple index on the dimension table's join column.

When SPD Server SQL is submitted that does not meet these STARJOIN conditions, SPD Server reverts to performing the requested SQL task using SPD Server's pairwise join strategy. The "[SPD Server STARJOIN Examples](#)" section of this document provides three examples that show valid, invalid, and restricted candidates for the SPD Server STARJOIN facility:

- [Example 1: Valid STARJOIN Candidate](#)
- [Example 2: Invalid STARJOIN Candidate](#)
- [Example 3: STARJOIN Candidate with Created or Calculated Columns](#)

Enabling STARJOIN Optimization in SPD Server

SPD Server STARJOIN optimization is enabled by default. The "[RESET Statement Options for SPD Server STARJOIN](#)" section provides detailed information on statement options that enable or disable the STARJOIN facility in SPD Server.

Invoking the SPD Server STARJOIN Facility

SPD Server knows when to use the STARJOIN facility because it is topology based. SPD Server invokes STARJOIN based on the SQL that is submitted. When SQL is submitted and STARJOIN optimization is enabled, SPD Server checks the submitted SQL for admissible STARJOIN patterns. SPD Server SQL identifies a fact table by scanning for a common equally joined table among multiple join predicates in a WHERE clause. When SPD Server SQL detects patterns that have multiple equally joined operators sharing a common table, the common table becomes the star schema's fact table.

When an SQL statement that is submitted to SPD Server uses structures that indicate the presence of a star schema, the STARJOIN validation checks begin.

SPD Server STARJOIN Optimization

The SPD Server STARJOIN optimization process searches for the most efficient SQL strategy to use for computations. The STARJOIN optimization process consists of three steps, regardless of the number of dimension tables that are joined to the fact table in the star schema.

1. Classify dimension tables that are called by SQL as Phase I tables or Phase II tables.
2. Phase I probes fact table indexes and selects a STARJOIN strategy.
3. Phase II performs index lookups and joins subsetted fact table rows with Phase II tables.

Classify Dimension Tables That Are Called by SQL as Phase I Tables or Phase II Tables

After the STARJOIN planner validates the join sub-tree, join optimization begins. Join optimization is the process that searches for the most efficient SQL strategy to use when joining the tables in the star schema.

The first step in optimization is to examine the dimension tables that were called by SQL for

structures that SPD Server can use to improve performance. Each dimension table is classified as a Phase I table or a Phase II table. The structure of a dimension table and whether the submitted SQL filters or subsets the table's contents determine its classification. SPD Server uses different processes to handle Phase I and Phase II dimension tables.

Phase I tables can improve performance. A Phase I table is a dimension table that is either very small (nine rows or less), or a dimension table whose SQL queries contain one or more filtering criteria that is expressed with a WHERE clause. A Phase II table is any dimension table that does not meet Phase I criteria. Rows in Phase II tables that are referenced in the SQL query are not subsetted.

Consider the star schema that is illustrated in [Figure 1](#), with the fact table Sales and the dimension tables Products, Supplier, Location, and Time.

Suppose a submitted SQL query requests transaction reports from the fact table Sales for all stores where the location is the state of North Carolina, for the time period of the month of January, for all products, and for all suppliers. The SQL query subsets the Location and Time tables, so SPD Server classifies the Location and Time tables as Phase I tables. The query requests information from all of the rows in the Product and Supplier tables. Because those tables are not subsetted by a WHERE clause in the submitted SQL, STARJOIN classifies the Products and Supplier tables in this query as Phase II tables.

Now, using the same star schema, add more detail to the SQL query. Set up a new query that requests transaction reports from the fact table Sales for all stores where the location is the state of North Carolina, for the time period of the month of January, and for products where the supplier is from the state of North Carolina. The subsetted dimension tables Location, Time, and Supplier are classified as Phase I tables. The Products table, unfiltered by the submitted SQL query, is classified as a Phase II table.

Dimension tables are classified as Phase I or Phase II tables because the two types of tables require different index probe methods.

Phase I Probes Fact Table Indexes and Selects a STARJOIN Strategy

Phase I uses the SQL join keys from the subsetted Phase I dimension tables to get a smaller set of candidate rows to query in the central fact table. After optimizing the candidate rows in the fact table, the Phase I index probe examines index structures to determine the best STARJOIN strategy to use. There are two SPD Server STARJOIN strategies: the IN-SET strategy and the COMPOSITE strategy. In all but a few cases, the IN-SET strategy is the most robust and efficient processing strategy. The user can determine which strategy SPD Server will choose by providing the required table index types in the submitted SQL.

Phase I creates the smaller set of candidate rows in the central fact table by eliminating fact

table rows that do not match the SQL join keys from the subsetted Phase I dimension tables. For example, if the SQL query requests information about transactions that occurred only in North Carolina store locations, the candidate rows that are retained in the fact table will use the SQL that subsets the Location dimension table:

```
WHERE location.STATE = "NC";
```

If the Sales fact table contains sales records for all 50 states, Phase I uses the SQL that subsets the Location dimension table to eliminate the sales records of all stores in states other than North Carolina from the fact table candidate rows. The example is simple, but powerful -- reducing the fact table candidate row set to transactions from only North Carolina stores eliminates massive amounts of nonproductive data processing.

The Phase I index probe inventories the number and types of indexes on the fact table and dimension tables as it attempts to identify the best STARJOIN strategy. To use the STARJOIN IN-SET strategy, Phase I must find simple indexes on all SQL join columns in the fact table and dimension tables. Otherwise, to use the STARJOIN COMPOSITE strategy, Phase I searches for the best composite index that is available on the fact table. The best composite index for the fact table is the composite index that spans the largest set of join predicates from the aggregated Phase I dimension tables.

Based on the fact table and dimension table index probe, SPD Server selects the STARJOIN strategy using the following logic:

- If one or more simple indexes are found on fact table and dimension table SQL join columns, and no spanning composite indexes are found on the fact table, SPD Server selects the STARJOIN IN-SET strategy.
- If an optimal spanning composite index is found on the fact table, and no simple indexes are found on fact table and dimension table SQL join columns, SPD Server selects the STARJOIN COMPOSITE strategy.
- If both simple and spanning composite indexes are found, SPD Server generally selects the STARJOIN IN-SET strategy, unless the composite index is an exact match for all of the Phase I join predicates, and only lesser matches are available with the IN-SET strategy.
- If no suitable indexes are found for either STARJOIN strategy, SPD Server does not use STARJOIN; it joins the sub-tree using the standard SPD Server pairwise join.

The IN-SET and COMPOSITE join strategies have some underlying differences.

The IN-SET join strategy uses an IN-SET transformation of dimension table metadata to produce a powerful compound WHERE clause to be used on the STARJOIN fact table. The "IN" part of the term "IN-SET" refers to an IN specification in the SQL WHERE clause. The

IN-SET is the set of values that populate the contents of the SQL IN query expression. For example, in the following SQL WHERE clause, the cities **Raleigh**, **Cary**, and **Clayton** are the values of the IN-SET:

```
WHERE location.CITY in ("Raleigh", "Cary", "Clayton");
```

For the IN-SET strategy, Phase I dimension tables are subsetted, and then the resulting set of join keys form the SQL IN expression for the fact table's corresponding join column. You must have simple indexes on all SQL join columns in both the fact table and dimension tables before STARJOIN Phase I can select the IN-SET strategy.

If the dimension table Location has six rows for Raleigh, Cary, and Clayton, then six STORE_NUMBER values are applied to the IN-SET WHERE clause that is used to select the candidate rows from the central fact table. The STARJOIN IN-SET facility transforms the dimension table's CITY values into STORE_NUMBER values that can be used to select candidate rows from the Sales fact table. The transformed WHERE clause to be applied to the fact table might resemble the following code:

```
WHERE fact.STORE_NUMBER in  
  (100,101,102,103,104,105,106);
```

You can use IN-SET transformations in a star schema that has any number of dimension tables and a fact table. Consider the following example dimension table subsetting statement:

```
WHERE location.CITY in  
  ("Raleigh","Cary","Clayton")  
and Time.SALES_WEEK = 1;
```

Because the Sales fact table has no matching CITY column to join with the Location dimension table, and no matching SALES_WEEK column to join with the Time table, the IN-SET strategy uses transformations to create a WHERE clause that the Sales fact table can resolve:

```
WHERE fact.STORE_NUMBER in  
  (100,101,102,103,104,105,106)  
and Time.SALES_DATE in  
  ('01JAN2005'd,'02JAN2005'd,'03JAN2005'd,  
  '04JAN2005'd,'05JAN2005'd,'06JAN2005'd,  
  '07JAN2005'd,);
```

The advantage of the STARJOIN facility is that it handles all of the transformations on a fact table, from dimension table subsetting to IN-SET WHERE clauses.

The COMPOSITE join strategy uses a composite index on the fact table to exhaustively probe

the full Cartesian product of the combined join keys that is produced by the aggregated dimension table subsetting. SPD Server compares the composite indexes on the fact table to the theoretical composite index that is made from all of the join keys in the Phase I dimension tables. Phase I selects the best composite index on the fact table, based on the join requirements of the dimension tables.

A disadvantage of using the COMPOSITE join strategy is that when more than a few join keys exist, the Cartesian product map can become large geometric matrixes that can interfere with processing performance. You must have a composite index on the fact table that consists of Phase I dimension table join columns before STARJOIN Phase I can select the COMPOSITE join strategy.

If any Phase I dimension tables contain join predicates that do not have supporting simple or composite indexes on the fact table, those Phase I dimension tables are dropped from Phase I processing and are moved to the Phase II group.

Phase II Performs Index Lookups and Joins Subsetted Fact Table Rows with Phase II Tables

Phase I optimizes the join strategies between the Phase I dimension tables and the candidate rows from the fact table . After Phase I terminates, Phase II takes over. Phase II completes the indicated joins between the candidate rows from the fact table and the corresponding rows in the subsetted Phase I dimension tables. After completing the joins with the Phase I dimension tables, Phase II performs index lookups from the fact table to the Phase dimension II tables. Phase II dimension tables should have indexes created on all columns that join with the fact table.

When SPD Server completes the STARJOIN Phase I and Phase II tasks, the STARJOIN optimizations have been performed, the STARJOIN strategy has been selected, and the subsetted dimension tables and fact table joins are ready to run and produce the desired SQL results set.

Indexing Strategies to Optimize STARJOIN Query Performance

Once the baseline criteria to create an SQL STARJOIN in SPD Server have been satisfied, you can configure indexing to influence which strategy the SPD Server STARJOIN facility chooses.

With the IN-SET strategy, the SPD Server STARJOIN facility can use multiple simple indexes on the fact table. The IN-SET strategy is the simplest to configure, and usually provides the best performance. To configure your work to choose the STARJOIN IN-SET strategy, create a simple index on each fact table and

dimension table SQL column that you want to use in a join relation. Creating simple indexes prevents STARJOIN Phase I from rejecting a Phase I dimension table so that it becomes a non-optimized Phase II table. In addition, simple indexes facilitate the Phase II fact-table-to-dimension-table join lookup.

Consider the following SQL code for a star schema with one fact table and two dimension tables:

```
PROC SQL;
select F.FID, D1.DKEY, D2.DKEY
from fact F, DIM1 D1, DIM2 D2
where D1.DKEY EQ F.D1KEY
and D2.DKEY EQ F.D2KEY
and D1.REGION EQ 'Midwest'
and D2.PRODUCT EQ 'TV';
```

[Indexing to Optimize the IN-SET Join Strategy](#)

The SPD Server IN-SET join strategy is the preferred strategy for almost every STARJOIN. If you want the previous example code to trigger the IN-SET STARJOIN strategy, create simple indexes on the join columns for the star schema's fact table and dimension tables:

- On the fact table F, create simple indexes on columns F.D1KEY and F.D2KEY.
- On the dimension tables D1 and D2, create simple indexes on columns D1.DKEY and D2.DKEY.

Other fact table and dimension table indexes might exist that could filter WHERE clauses, but those simple indexes are the indexes that are needed to enable the STARJOIN IN-SET join strategy.

[Indexing to Optimize the COMPOSITE Join Strategy](#)

Using the COMPOSITE join strategy, the dimension tables with WHERE clause subsetting are collected from the set of equally joined predicates. A fact table composite index is needed for the fact table columns that correspond to the subsetting dimension table columns. The composite index on the fact table is necessary to facilitate the dimension tables' Cartesian product probes on the fact table rows. The STARJOIN optimizer code looks for the best composite index, which is based on the best and simplest left-to-right match of the columns in the COMPOSITE join.

If the subsetting in a STARJOIN is limited to a single dimension table, then the COMPOSITE join strategy can be enabled by creating a simple index on the join column of the single dimension table. That index is used to perform the Phase II index lookup on the fact table candidate rows. The fact table candidate row set is the result of the Phase I composite index probe.

For the previous example code to trigger the COMPOSITE STARJOIN strategy, create a composite index named COMP1 on the fact table for each of the dimension table keys: F. COMP1=(D1KEY D2KEY).

Other fact table and dimension table indexes might exist that could filter WHERE clauses, but the COMPOSITE index named COMP1 is the type of index that is needed to enable the STARJOIN COMPOSITE join strategy.

Although the COMPOSITE join strategy might appear to be a simpler configuration, the strongest utility of the COMPOSITE join strategy is limited to join relations between the fact table and dimension tables that are based on a Cartesian matrix of outcomes. As the number of dimension tables and join relations increases, the resulting Cartesian matrixes increase geometrically in size and can become unmanageable. The superior performance of the IN-SET strategy is so dramatic and robust that you should consider using the COMPOSITE join strategy only if you have good evidence that it compares favorably with the IN-SET strategy.

Example: Indexing using the IN-SET Join Strategy

The example star schema in [Figure 1](#) has four dimension tables (Supplier, Products, Location, and Time) and one fact table (Sales) with simple indexes on the SUPPLIER_ID, PRODUCT_CODE, STORE_NUMBER, and SALES_DATE columns in the Sales fact table.

Consider the following SQL query to create a January sales report for an organization's North Carolina stores:

```
PROC SQL;
select
  sum(s.sales_amt) as sales_amt
  sum(s.units_sold) as units_sold
  s.product)code,
  t.sales_month

from
  spdslib.sales s,
  spdslib.supplier sup,
  spdslib.products p,
  spdslib.location l,
  spdslib.time t

where
  s.store_number = l.store_number
and s.sales_date = t.sales_date
and s.product_code = p.product_code
```



```

and s.supplier_id = sup.supplier_id
and l.state = 'NC'
and t.sales_date
    between '01JAN2005'd and '31JAN2005'd;

quit;

```

During optimization, the STARJOIN planner examines the WHERE clause subsetting in the SQL to determine which dimension tables qualify as Phase I tables and which are Phase II tables.

The WHERE clause subsetting of the STATE column of the Location dimension table (`where ... l.state = 'NC'`) and the subsetting of the SALES_DATE column of the Time dimension table (`where ... t.sales_date between '01JAN2005'd and '31JAN2005'd`) cause SPD Server to process the Location and Time tables as Phase I tables. The remaining dimension tables Supplier and Products are processed as Phase II tables.

SPD Server STARJOIN uses the Phase I dimension tables to reduce the rows in the fact table to candidate rows that contain the matching criteria. The values in each dimension table key are used to create a list of values that meet the subsetting criteria of the fact table.

For example, the previous SQL query is intended to create a January sales report for stores located in North Carolina. Note that the WHERE clause in the SQL code joins the Location and Sales tables on the STORE_NUMBER column. Suppose that there are 10 unique North Carolina stores, with consecutively ordered STORE_NUMBER values that run from 101 to 110. When the WHERE clause is evaluated, the results will include a list of the 10 North Carolina store IDs that existed in January 2005.

With simple indexes on the fact table and dimension tables for the STORE_NUMBER column, STARJOIN chooses the IN-SET strategy. Subsetting the STATE column values to 'NC' allows STARJOIN to build the set of store numbers that are associated with North Carolina locations. STARJOIN can use the set of North Carolina store numbers to generate a `where ... in` SQL expression. SQL uses the `where ... in` expression to efficiently subset the candidate rows in the fact table before the final SQL expression evaluation.

In other words, STARJOIN uses a matrix of database relationships and index combinations to reorganize the SQL expression for more internal processing that can take advantage of the IN-SET join strategy. For the previous example code, the internal STARJOIN SQL reorganization resembles the following example code. The WHERE clause IN-SET statements for the STORE_NUMBER and TIME columns can be rapidly processed to subset the candidate rows in the Sales fact table. (The optimized code sections are highlighted.)

```

PROC SQL;
select

```

```

sum(s.sales_amt) as sales_amt
sum(s.units_sold) as units_sold
s.product)code,
t.sales_month

from
  spdslib.sales s,
  spdslib.supplier sup,
  spdslib.products p,
  spdslib.location l,
  spdslib.time t

where
  s.store_number = l.store_number
and s.sales_date = t.sales_date
and s.product_code = p.product_code
and s.supplier_id = sup.supplier_id
and s.store_number
  in (101,102,103,104,105,106,107,108,109,110)
and s.time
  in ('01JAN2005'd,'02JAN2005'd,'03JAN2005'd,
      '04JAN2005'd,      <...>, '28JAN2005'd,
      '29JAN2005'd,'30JAN2005'd,'31JAN2005'd);

quit;

```

After Phase I completes the candidate row optimization on the Sales fact table, Phase II processes the optimized query from the fact table outward. Phase II uses the values in the fact table's subsetted candidate rows to perform index lookups on the dimension tables' contents to complete the join in the most efficient manner.

SPD Server STARJOIN RESET Statement Options

SPD Server recognizes several RESET statements that can configure or provide information about the STARJOIN facility in SPD Server SQL.

- [RESET NOSTARJOIN](#)
- [RESET STARMAGIC](#)
- [RESET DETAILS = "stj\\$"](#)
- [Example: NOSTARJOIN, STARMAGIC, and stj\\$](#)

[RESET NOSTARJOIN=\[0/1\]](#)

The NOSTARJOIN option suppresses the use of the SPD Server STARJOIN optimizer in the planning and running of SQL statements that have valid STARJOIN patterns or star schemas. The statements NOSTARJOIN and NOSTARJOIN=1 are equivalent. When NOSTARJOIN is enabled, SPD Server ignores STARJOIN and uses pairwise joins to plan and run SQL statements. The default setting is NOSTARJOIN=0, meaning that in SPD Server, STARJOIN is enabled unless reset, and STARJOIN optimization occurs when SQL recognizes a valid SPD Server pattern or star schema.

RESET STARMAGIC=*nnn*

STARMAGIC is the STARJOIN counterpart to the SQL MAGIC number option. You can set magic numbers that direct STARJOIN to override internal heuristics, which results in enhanced join strategies. The STARMAGIC option uses bit flags to configure the STARJOIN code. You can select different controls by adding the values for the different bit flags in the following STARMAGIC set:

Value	Meaning
1	forces all dimension tables to be classified as Phase I tables.
2	obsolete; not used.
4	requires exact matches on the fact table composite index to meet STARJOIN Phase I conditions.
8	disables IN-SET join strategy. (Default setting is enabled.)
16	disables COMPOSITE join strategy. (Default setting is enabled.)

RESET DETAILS="stj\$"

All internal STARJOIN debugging information is tied to RESET DETAILS="stj\$". Issuing this statement displays available information as SPD Server attempts to validate a join sub-tree. The RESET DETAILS="stj\$" statement is useful for debugging STARJOIN and SQL statements.

Example: STARJOIN RESET Statements

The following example connects to sasspds, and then issues the "stj\$" RESET option to display all available information as SPD Server attempts to validate the join sub-tree for the submitted SQL on a star schema. The STARMAGIC=16 setting disables the STARJOIN COMPOSITE join strategy (STARJOIN COMPOSITE joins are enabled by default in SPD Server). The NOSTARJOIN=0 setting means that STARJOIN is enabled (or resets a disabled STARJOIN facility) and ensures that STARJOIN optimization occurs if SQL recognizes a valid SPD Server pattern or star schema. (The STARJOIN facility is enabled by default in SPD Server.)

After submitting the following SQL statements, the code disconnects from sasspds and quits:

```
PROC SQL;
  connect to sasspds
    (dbq="star"
     server=sunburn.5007
     user='anonymous' );

  execute (reset
    DETAILS="stj$"
    STARMAGIC=16
    NOSTARJOIN=0)

  by sasspds;

  execute (
    ...
    SQL statements
    ...);
  by sasspds;

  disconnect from sasspds;
quit;
```

SPD Server STARJOIN Examples

Example 1: Valid SQL STARJOIN Candidate

The following code is an example of an SQL submission that SPD Server is able to use as a star schema. The submission is a valid candidate because:

- a single central fact table, Sales, exists
- the dimension tables Time, Products, Location, and Supplier all join with the fact table Sales
- each dimension table appears in only one join condition
- all dimension tables link to the fact table using equally joined operators

```

PROC SQL;
create table Sales_Report as
select a.STORE_NUMBER,
       b.quarter
       c.month,
       d.state,
       e.SUPPLIER_ID

sum(a.total_sold) as tot_qtr_mth_sales
from   Sales a,
       Time b,
       Products c,
       Location d,
       Supplier e

where a.sales_date   = b.sales_date
      and a.STORE_NUMBER = d.store_number
      and a.PRODUCT_CODE = c.product_code
      and a.SUPPLIER_ID = d.supplier_id
      and b.quarter in (3, 4)
      and c.PRODUCT_CODE in (23, 100)

group by b.quarter,
         a.STORE_NUMBER,
         b.month;
quit;

```

[Example 2: Invalid SQL STARJOIN Candidate](#)

The following code is an example of an SQL submission that SPD Server is not able to use as a star schema because no single central fact table can be identified. Changes to the previous code example are highlighted:

```

PROC SQL;
create table Sales_Report as
select a.STORE_NUMBER,
       b.quarter

```

```

        c.month,
        d.state,
        e.SUPPLIER_ID

sum(a.total_sold) as tot_qtr_mth_sales
from   Sales a,
       Time b,
       Products c,
       Location d,
       Supplier e

where a.sales_date    = b.sales_date
      and a.STORE_NUMBER = d.store_number
      and a.PRODUCT_CODE = c.product_code
      and c.SUPPLIER_ID = d.supplier_id
      and b.quarter in (3, 4)
      and c.PRODUCT_CODE in (23, 100)

group by b.quarter,
         a.STORE_NUMBER,
         b.month;

quit;

```

SPD Server is not able to use the SQL submission in this example as a star schema. This submitted code joins the dimension tables for Time, Products, and Location to the Sales table, but the table for Supplier is joined to the Sales table through the Products table. As a result, the topology does not define a single central fact table.

Example 3: STARJOIN Candidate with Created or Calculated Columns

The STARJOIN facility in SPD Server supports calculated or created columns. The following code is an example of an SQL submission that creates columns, but still uses STARJOIN optimization, if the central fact table and the dimension tables contain indexes on the join columns for the STARJOIN:

```

PROC SQL;
create table &Regional_Report as
select case d.state
       when 1 then 'NC'
       when 2 then 'SC'
       when 3 then 'GA'
       when 4 then 'VA'
       else ' '

```

```

end as state_abv,
b.quarter,
sum (a.tot_amt) as total_amt

from wk_str_upd_t a,
     week_t b,
     location_t d,

where a.we_dt      = b.we_dt
     and a.chn_str_nbr = d.chn_str_nbr
     and b.quarter  = 2

group by d.state,
         b.quarter
having d.state in (1,2,3,4);
quit;

```

The highlighted code creates a column called `state_abv`. The SPD Server STARJOIN facility supports created columns if the appropriate indexes on the join columns exist in the fact table and dimension tables.

Copyright © 2008 by SAS Institute Inc., Cary, NC, USA. All rights reserved.

SAS Scalable Performance Data (SPD) Server SQL Query Rewrite Facility

Contents

- [Overview of the SQL Query Rewrite Facility](#)
- [Configuring Storage Space for the SQL Query Rewrite Facility](#)
- [SQL Query Rewrite Facility Options](#)
 - [_QRWENABLE Option](#)
 - [Examples](#)
 - [_QRW Option](#)
 - [Examples](#)

[Overview of the SQL Query Rewrite Facility](#)

The SQL Query Rewrite facility in SPD Server examines SQL queries in order to optimize processing performance. Some SQL queries contain SQL statements and sub-queries that can be more rapidly evaluated in a separate space, as opposed to sequentially evaluating large blocks of SQL statements. When SPD Server detects and processes SQL statements or sub-queries in a separate space, intermediate result tables are produced. The original SQL query is dynamically rewritten, using intermediate results tables to replace the SQL code that was separately evaluated. The result is a dynamic process that evaluates and processes SQL queries more efficiently.

Inserting the derived intermediate data into the original SQL query does not change the quantitative results; it only expedites the processing that is required to calculate them. The SQL Query Rewrite Facility does not change the content of the query's answer row set. However, the *order* of the rows in the query answer set may vary if you compare the optimized query answer set with the query answer set SPD Server generates with the SQL Query Rewrite facility disabled.

[Configuring Storage Space for the SQL Query Rewrite Facility](#)

The SQL Query Rewrite Facility uses intermediate results tables to expedite original SQL queries. SPD Server administrators must provide adequate space for the generation and storage of the

intermediate results tables. The intermediate results tables are formatted as SPD Server tables. Optional indexes are permitted.

Intermediate results tables are stored in a common SPD Server LIBNAME domain that the SPD Server administrator specifies. One dedicated SQL Rewrite Facility LIBNAME domain is sufficient to provide adequate intermediate results table storage for many concurrent SPD Server users. Why is only one domain enough? The SQL Query Rewrite Facility utilizes the SPD Server TEMP=YES option setting when accessing the LIBNAME domain for intermediate result tables. The TEMP=YES option creates a processing environment where multiple SPD Server users can concurrently create tables with no name or resource contention issues. The TEMP=YES option also automatically cleans up the contents of the working storage area when users close their SPD Server session in a normal fashion.

SPD Server administrators and users can both specify LIBNAME domains for SQL Query Rewrite Facility intermediate results storage. SPD Server administrators can configure use the TMPDOMAIN= parameter in the spdsserv.parm file to specify the SQL Query Rewrite Facility intermediate results storage domain:

```
TMPDOMAIN=<DomainName>;
```

where *<DomainName>* is the name of a LIBNAME domain that is defined in the SPD Servers associated libnames.parm file.

SPD Server users can override the primary TMPDOMAIN= location by specifying their own LIBNAME domain for SQL Query Rewrite Facility intermediate results storage. Users specify their own LIBNAME domain by using the pass-through SQL RESET command with the TMPDOMAIN= option. For example, if an individual SPD Server user wanted to use the EMATMP LIBNAME domain for SQL Rewrite Facility intermediate results, the user would submit the following RESET command in his or her SQL job stream:

```
execute(reset tmpdomain=ematmp) by sasspds;
```

Setting TMPDOMAIN=EMATMP causes the EMATMP domain to take precedence over the TMPDOMAIN= setting that was specified in the spdsserv.parm file. Any LIBNAME domain that that an individual user submits as an SQL Query Rewrite storage location must be defined in the libnames.parm file of the SPD server that runs the pass-through SQL code.

Reassigning the SQL Query Rewrite Facility intermediate results storage location does not affect TEMP=YES environment setting that permits concurrent access to tables in the domain by multiple SPD Server users. This means that multiple SPD Server users can specify and share remapped TMPDOMAIN= locations without table handling or contention issues.

Note: If the SPD Server parameter TMPDOMAIN is not configured and the SQL query rewrite is enabled, the query rewrite will fail with the following error:

SPDS_ERROR: Error materializing RWE context.

SQL Query Rewrite Facility Options

The SQL Query Rewrite Facility is enabled by default in SPD Server. That means when an SPD Server user submits SQL statements that contain sub-expressions that SPD Server can handle more efficiently by using the SQL Query Rewrite Facility, the software will optimize the SQL query. RESET options provide control over the SQL Query Rewrite Facility.

- [_QRWENABLE Option:](#)
- [_QRW Option](#)

_QRWENABLE Option: Use the `_QRWENABLE` reset option to disable the SQL Query Rewrite Facility. You might use this option if you suspect that the SQL Query Rewrite Facility is not enhancing the performance of the SQL query execution. The SQL Query Rewrite Facility is enabled by default.

Examples:

Disable SQL Query Rewrite Facility:

```
execute(reset no_qrwenable) by sasspds; /* Disable
query rewrite */
execute(reset _qrwenable=0) by sasspds; /* Another
way to disable */
```

Re-enable SQL Query Rewrite Facility:

```
execute(reset _qrwenable) by sasspds; /* Re-enable
query rewrite */
execute(reset _qrwenable=1) by sasspds; /* Another
way to enable */
```

_QRW Option: Use the _QRW reset option to enable diagnostic debugging and tracing outputs from the SQL Query Rewrite Facility in the log. The diagnostic debugging option is disabled by default.

Examples:

Enable diagnostic debugging function:

```
execute(reset _qrw) by sasspds; /* Enable
diagnostics */
execute(reset _qrw=1) by sasspds; /* Another way to
enable */
```

Disable diagnostic debugging function:

```
execute(reset no_qrw) by sasspds; /* Disable
diagnostics */
execute(reset _qrw=0) by sasspds; /* Another way to
disable */
```

Copyright © 2008 by SAS Institute Inc., Cary, NC, USA. All rights reserved.