



SAS/CONNECT® Software

Maximize computing resources and manage distributed applications

What does SAS/CONNECT® software do?

SAS/CONNECT software establishes connections between networked computers with various operating systems and enables organizations to combine the resources of these machines for maximum benefit. It efficiently distributes computing workloads across different CPUs and provides scalability through parallel processes.

Why is SAS/CONNECT® software important?

With the ability to access, manage and process data in distributed environments, you can maximize use of all computing resources to shorten execution times and speed decision making.

For whom is SAS/CONNECT® software designed?

This product is designed for those who need to more fully exploit their distributed computing resources, including SMP hardware, and provide users with access to data and increased processing power from a multitude of systems.

Complex infrastructures are a way of life for corporate IT organizations – a problem exacerbated by mergers and acquisitions. New data sources, a conglomeration of hardware and increasing numbers of users demanding accurate information – is there any way to maximize use of your diverse resources and deliver information more quickly?

To achieve cost effectiveness and operating efficiency, IT organizations must optimally harness system resources and fully leverage previous outlays. SAS/CONNECT software easily establishes connections between diverse networked computers and enables you to exploit all CPUs for maximum benefit.

SAS/CONNECT software provides the essential tools for sharing data, applications and processing power across multiple computing environments. It gives you control over where and how to execute each part of an application and takes advantage of parallel, multi-threaded processing to maximize use of hardware resources.

By dividing time-consuming tasks into multiple units of work and executing these units in parallel, jobs run in substantially less time than if each task is performed sequentially. In addition, you can easily move SAS® files and flat files across different machine architectures and SAS releases. Data translations and release conversions take place automatically.

With SAS/CONNECT software's cross-platform data access and parallel-processing capabilities, your programmers can quickly assimilate mountains

of distributed data, speed processing time for large jobs and dramatically reduce the time it takes to get business information to those who need it.

Key Benefits

- Reduce costs by fully exploiting existing IT infrastructures.** SAS/CONNECT software enables you to integrate hardware, applications and data from distributed, heterogeneous platforms and provides an easy way to minimize idle CPU cycles. This reduces the need for constant hardware upgrades to meet growing CPU, memory and I/O requirements. SAS/CONNECT provides a code-efficient way to distribute applications across diverse platforms and helps you maximize use of computing resources 24/7.
- Speed up application performance and lessen the time it takes to transform and analyze data.** You can divide jobs into independent units and distribute them across multiple CPUs in your environment.
- Eliminate data silos and isolated islands of information.** With SAS/CONNECT software, you can integrate and access data and applications on remote systems throughout the enterprise. SAS offers a unique multiplatform environment as well as transparent representation of data from different platforms – eliminating manual encoding and decoding of data from various systems. Make use of all data and resources to provide decision makers with a holistic view of the organization.



Product Overview

SAS/CONNECT software makes it easy to distribute work to multiple CPUs across platforms and operating systems. You can move the application to where the data is or move the data to where the application executes, whichever delivers results more quickly and efficiently.

In addition, SAS/CONNECT software can receive input from third-party schedulers, allowing you to maximize computing resources by executing applications on the most appropriate platforms 24/7, particularly at off-peak hours.

With SAS/CONNECT software, SAS applications can run between releases as well as across heterogeneous platforms. MultiVendor Architecture™ enables the software to run as a virtual machine so programmers can develop and test jobs on the system of their

choice and then easily distribute the jobs in complex operating environments to speed up execution. SAS/CONNECT interoperates with all SAS components, including the SAS Scalable Performance Data Engine, SAS Scalable Performance Data Server® and threaded SAS procedures, which further speed processing and I/O throughput by incorporating parallel threading. In addition, SAS/CONNECT allows Java clients to access and exploit the power of a SAS server.

Key Capabilities

Compute services

SAS/CONNECT software's compute services enable the use of multiprocessor and remote computing resources, including hardware, software and data, to most efficiently execute an application. You can move any or all portions of an application's processing to other processors (either local or remote) to take

advantage of hardware resources, use software in remote environments, interface with legacy systems and execute against a remote copy of data.

The results of the remote processing are then returned to the local machine. This is useful when the remote machine has hardware and/or software that can more efficiently perform a particular task. Compute services are also helpful if the amount of data to be processed is too large to move to the local machine or if the data is updated too frequently for a local static copy to be useful.


Data transfer services

SAS/CONNECT software's data transfer services provide a method for moving a copy of data from one machine to another, where a physical copy is then created. Subsequent local processing takes place against the local copy of the data without generating further network traffic until you decide to update the copy with another transfer.

Data transfer services automatically perform any conversion or translation necessary to move data, such as from one SAS release to another or from one machine representation to another. In addition, these services can move data stored in SAS data sets, external databases, SAS catalogs, SQL views and external files.

Remote library services

SAS/CONNECT software's remote library services provide access to remote data libraries as if they were stored locally. The data moves through the network only as the local execution requests it. A copy of the data is not written to the local file system, and the data must pass through the network on subsequent use by the local processing. This allows you to



Move your application to where the data is or move data to where the application is, whichever more efficiently delivers results. With SAS/CONNECT software, you can access and transfer remote data as well as partition and execute applications across processors – enabling you to maximize all of your computing resources.

maintain a single copy of your data and build applications that provide seemingly identical access to local and remote data without requiring the user to know where the data resides.

Independent parallelism

SAS/CONNECT includes a multiprocessing feature (MP CONNECT) that provides independent parallelism so you can divide your application into independent pieces that can be executed simultaneously and then merged into a final result set. This type of asynchronous, parallel processing can occur on the same machine or on different machines. If you have symmetric multiprocessing (SMP) machines in-house, you can immediately benefit from faster results. Additionally, you can leverage unused processor capacity across your network.

Pipeline parallelism

The piping feature of SAS/CONNECT speeds throughput by allowing data to be piped from one process to another between SMP or distributed processors. Pipeline parallelism allows overlapped execution of SAS DATA steps and/or SAS procedures, which require only a single pass through the data. Rather than having to wait for one process to completely finish writing output before another begins, piping enables the waiting process to begin as soon as the first process starts generating data.

In addition, piping the data eliminates the intermediate write to disk, saving both time and disk space. SAS/CONNECT software enables pipeline processing on SMP hardware as well as on machines across a network.

Key Features

Compute services

- Enables you to distribute processing to the most appropriate platforms in your network.
- Enables asynchronous, parallel processing on SMP hardware as well as any number of platforms across your network.
- Minimizes total elapsed execution time.
- Bridges heterogeneous platforms, disparate data and software resources across your enterprise.

Data transfer services

- Enables you to easily copy SAS files, SAS programs and raw data files across the network.
- Automatically translates and converts data from one platform to another.
- Reduces time for large data transfers with file compression algorithms.
- Supports encryption of data as it flows on the network.

Remote compute services

- Lets you seamlessly access remote server data libraries as if they were stored locally.
- Provides GUI applications executing in the client session with read and write access to remote data.

Independent parallelism

- Segment applications so the appropriate pieces execute on the desired platform.
- Make efficient use of all hardware, including SMP machines.
- Tap unused processor cycles across any number of heterogeneous machines in your enterprise.

Pipeline parallelism

- Overlap processing execution by piping output from one process as the input into the next process in the pipeline.
- Improve performance.
- Reduce the demand for disk space.

Libref inheritance

- Eliminates the need to duplicate data for use in multiple sessions.
- Allows libraries that are defined in the client session to be inherited by multiple server sessions for read and write access.
- Child sessions can transparently access the parent data library for input parameters.
- Child sessions can use the parent library to store computations for easier consolidation.

Result distribution and sharing

- Distribute results by publishing to multiple channels, including the SAS BI report/content repository, the SAS Stored Process Server, printers, Microsoft Office documents, and email.
- Update analyses and results on a periodic basis using the native Windows scheduler.
- Export results to other applications such as Adobe Acrobat, Microsoft Excel, Microsoft Access or Microsoft Word.
- Surface results within a report or publish via stored services (also known as SAS Stored Processes) for use within other SAS applications such as the SAS Add-In for Microsoft Office and SAS Web Report Studio or other third-party clients.

SAS/CONNECT® Software System Requirements

To learn more about SAS/CONNECT system requirements, download white papers, view screenshots and see other related material, please visit www.sas.com/connect.



SAS Institute Inc. World Headquarters +1 919 677 8000

To contact your local SAS office, please visit: **www.sas.com/offices**

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 trademarks of their respective companies.
Copyright © 2011, SAS Institute Inc. All rights reserved. 101229_S74878.0711