

So Much Software, So Little Time: Deploying SAS® onto Oodles of Machines

Mark Schneider, SAS Institute Inc., Cary, NC

ABSTRACT

Distributing SAS® software to a large number of machines can be challenging at best and exhausting at worst. Common areas of concern for installers are silent automation, network traffic, ease of setup, standardized configurations, maintainability, and simply the sheer amount of time it takes to make the software available to end users. We describe a variety of techniques to ease the pain of provisioning SAS software, including the new stand-alone SAS® Enterprise Guide® and SAS® Add-In for Microsoft Office installers, as well as the tried and true SAS® Deployment Wizard record and playback functionality. We also cover ways to shrink SAS Software Depots, like the new "subsetting recipe" feature, in order to ease scenarios requiring depot redistribution. Finally, we touch on alternate methods for workstation access to SAS client software, including application streaming, desktop virtualization, and Java Web Start.

INTRODUCTION

The process of distributing software from a central location to multiple machines in your environment is known as "provisioning." Throughout this paper, we will refer to the central location from which you provision software as the "source machine" and computers to which you provision the software as the "target machines."

There are many provisioning techniques you can use when deploying the same set of SAS software to multiple end users. Some approaches, like the SAS Deployment Wizard record and playback functionality, involve running installs on all target machines. Other approaches, like Java Web Start, don't involve any deployment activity on the target machine besides installing supporting third-party software (e.g., Java).

Our intent is not to provide exhaustive technical detail on each of the approaches, but rather to offer a survey of options available. You will likely find value in employing a combination of the covered techniques, rather than choosing just one. Advantages and disadvantages are provided for each option to help you choose the right approach for each software distribution challenge you face.

RECORD AND PLAYBACK

The most flexible and robust way for deploying SAS software in unattended fashion is using the "Record and Playback" feature of the SAS Deployment Wizard. With this method you first stage an initial machine, deploying the desired set of SAS software via the SAS Deployment Wizard GUI. While doing so, you direct the SAS Deployment Wizard to capture all answers to dialog boxes in a response file. You can then feed this response file into subsequent SAS Deployment Wizard executions on other machines to deploy the exact same set of software in the same way. Note that you don't have to install SAS software on the source machine. All you need to do is capture SAS Deployment Wizard responses. You can leave actual installation to the target machines.

If you've placed your SAS Software Depot on a network drive accessible to all target machines, and your network bandwidth can accommodate the transfer of gigabytes of data to each during the deployment, this can be the easiest option. If you do not have such a network drive, but have sufficient network bandwidth to accommodate redistributing your SAS Software Depot to all target machines, this is a viable option as well. You may want to explore the depot subsetting feature of the SAS Deployment Wizard if you choose this route.

To insulate end users from needing to know about the deployment process, you are advised to create an operating system script that invokes the SAS Deployment Wizard out of the SAS Software Depot, and specifies the `-quiet`, `-wait`, and `-responsefile` options. You can then provision this script, along with the response file (and SAS Software Depot if necessary) to execute remotely on all target machines.

If you want end users to be able to initiate their own deployments and modify your recorded responses to dialog boxes, you can provide them with a script that omits the aforementioned `-quiet` option. All dialog boxes will be seeded with the responses you captured previously, but can be overridden by the end user. If you want end users to be able to specify some, but not all, of the dialog box responses, you can specify the `-partialprompt` option. This will hide any dialog boxes for which there is a response in the response file, but prompt for all others.

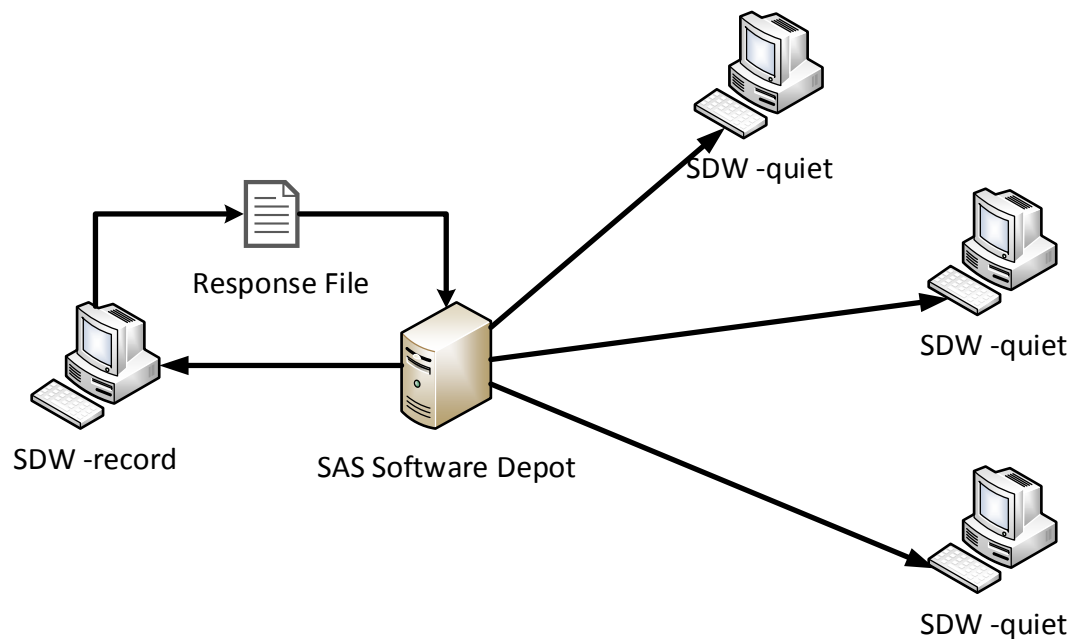


Figure 1. Record and Playback

Anything that can be deployed interactively through the SAS Deployment Wizard can be redeployed on other machines using this method. See the “Record and Playback” section of the [SAS Deployment Wizard and SAS Deployment Manager 9.4: User's Guide](#) for more details. Furthermore, all administrative tasks available in the SAS Deployment Manager may be recorded and played back in similar fashion.

Advantages:

- Supports all deployments, including installation and configuration
- Integrates with existing provisioning software like Microsoft System Center Configuration Manager (SCCM) and Tivoli
- Includes and installs third-party system requirements (e.g., .NET Framework) as needed

Disadvantages:

- Requires target machine access to a SAS Software Depot
- Heavyweight and frequently slower than other methods

STAND-ALONE INSTALLERS

A limited number of SAS products can be deployed using stand-alone installers delivered in your SAS Software Depot. In the `standalone_installs` directory at the root of your depot, you will find all such installers applicable to your software order. If you drill down into a specific product and version under the `standalone_installs` directory, you will find a README file as well as a single executable file. The executable file contains everything needed to install the product, including most third-party system requirements (.NET Framework is the noted exception due to its size). If you intend to install only SAS products having stand-alone installers on your target machines, and you do not anticipate installing any other SAS software to these machines in the future, this is likely your easiest option.

The two most common SAS desktop clients, SAS Enterprise Guide and SAS Add-In for Microsoft Office, are both supported in this fashion. The stand-alone installers for these products support both 32- and 64-bit environments. In each case, the same executable can be used on all target Windows machines, assuming they meet minimum system requirements as outlined at <http://support.sas.com/documentation/installcenter/94/win/index.html>.

If maintenance is available for a product, it will be included in its stand-alone installer. If the product was previously installed on the target machine, the maintenance will be applied. If the product was not previously installed, the original content plus maintenance will be installed. In all cases, the stand-alone installer determines the correct course of action and does it.

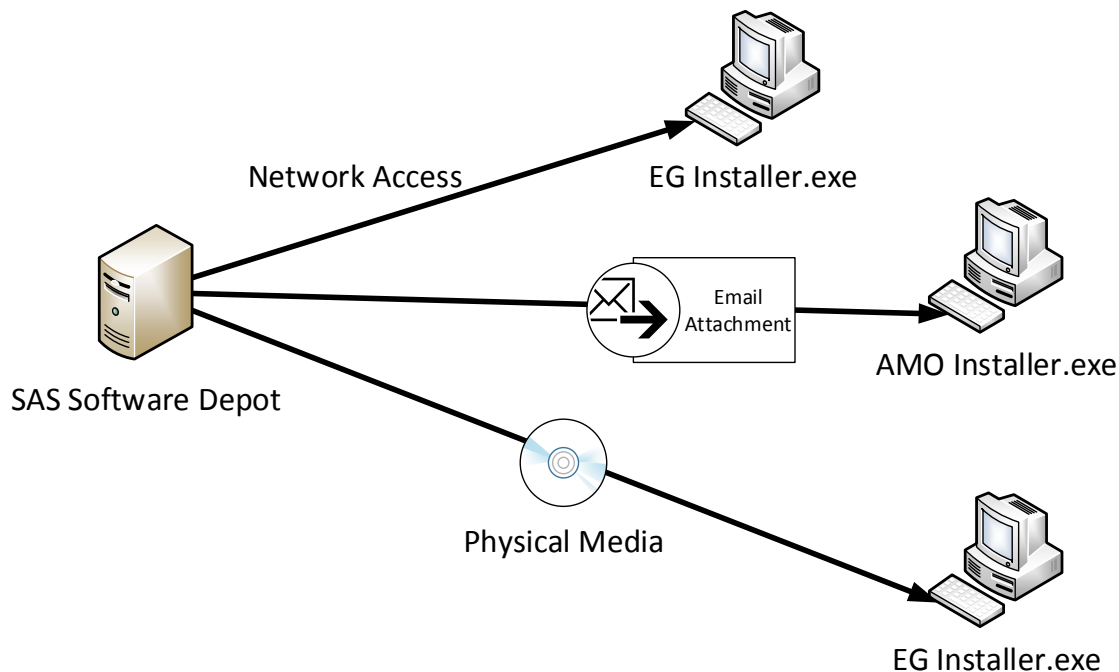


Figure 2. Stand-alone Installers

As with the SAS Deployment Wizard, all stand-alone installers support unattended mode. The same options (`-quiet`, `-responsefile`, `-wait`, and `-partialprompt`) are used, in addition to the `/verysilent` option, which is unique to the stand-alone installers. As such, these installers are well-suited for third-party provisioning software like SCCM and Tivoli. See the [SAS Deployment Wizard and SAS Deployment Manager 9.4: User's Guide](#) for more details.

Advantages:

- Lightweight
- Minimal dialog boxes
- Doesn't require target machine access to a SAS Software Depot
- Easy distribution via network, web download, e-mail, or physical media

Disadvantages:

- Limited product set
- Cannot mix stand-alone installers and SAS Deployment Wizard installs on the same machine
- Available only on Windows
- Must preinstall Microsoft .NET Framework separately

DELTA IMAGES

In some situations, you can use third-party software to capture delta images containing any changes made to a system as a result of a SAS deployment. These delta images can then be distributed to target machines that are updated in place with the changes.

Your first step is to create a virtual image representing your target machines. You take a snapshot of this virtual image, and then deploy your desired SAS software to the image. The third-party software compares the initial snapshot with the post-deployment image and records all changes in the environment, such as new or modified files and Windows registry keys. The changes are packaged up and distributed to target systems where they are unpackaged and deployed. On Windows, third-party software like Flexera's AdminStudio might create a Windows Installer (MSI) file to be distributed to the systems in your environment via a provisioning tool like Microsoft System Center Configuration Manager.

Other third-party software packages use the Windows Image (WIM) format to redistribute delta packages. One example of such software is SmartDeploy. Not only does it support this simplified delta image model, but it allows you to use the same delta packages for multiple versions of Windows operating systems. It accomplishes this by isolating operating system-specific content (e.g., device drivers) into "platform packs." The operating system-independent content lives in separate space.

Microsoft's ImageX and Deployment Image Servicing and Management (DISM) tools offer similar functionality; however, they cannot be used to produce delta images. They simply allow you to maintain full versions of operating systems in WIM format, and to redistribute those images in full. You can't pick and choose specific parts to be redistributed from an ImageX or DISM image.

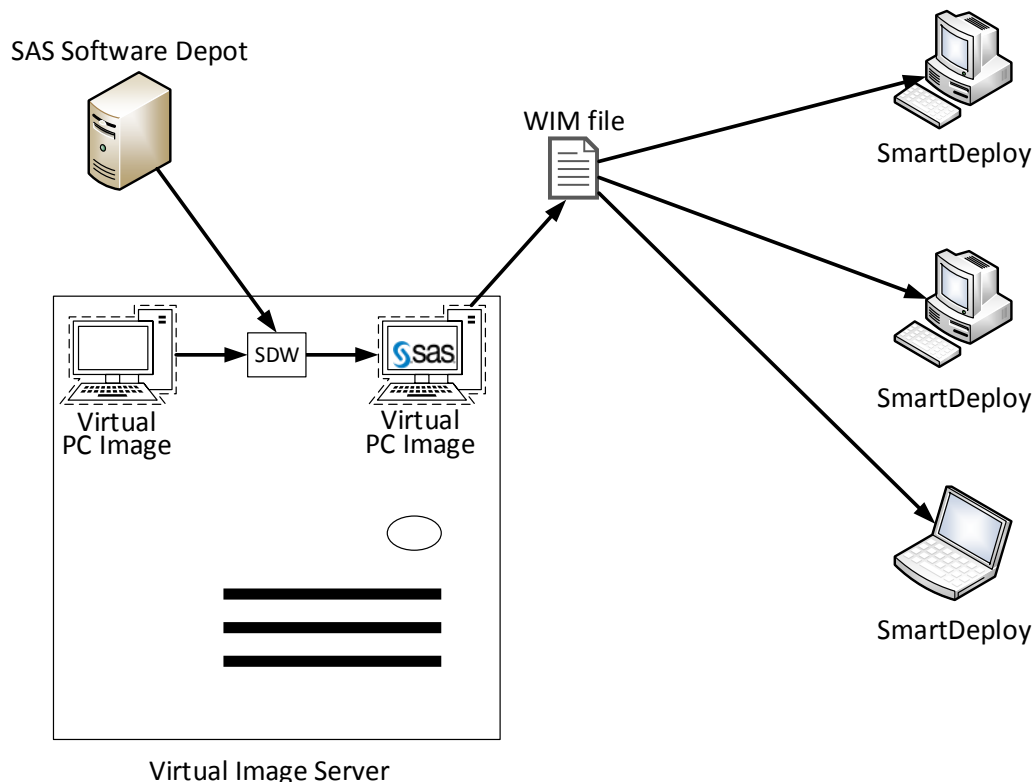


Figure 3. Delta Images

Advantages:

- Locally installed images that are typically faster
- No reliance on network access at SAS runtime allowing for disconnected operation
- Accommodates differences in target computers, including OS versions

Disadvantages

- Requires third-party image management software
- Is available only for Windows

WEB DOWNLOADS

Some SAS software is available directly from <http://support.sas.com/demosdownloads/setupintro.jsp>. Software at this site is typically intended for installation on client machines on which no other SAS software is installed. Much of this software is included with regular orders and can be installed from a SAS Software Depot using the SAS Deployment Wizard, including using the aforementioned Record and Playback functionality. However, there are times when users don't have access to a depot, and simply want to download client software that doesn't require a license and can be installed on the target machine or machines in isolation from any other SAS software.

An example product that is available for web download is SAS® Universal Viewer. This product allows users to view, sort, and filter SAS data sets even if they don't have SAS installed on their local machine. Installing products like SAS Universal Viewer involves going to the aforementioned website, reading the provided README file associated with the product, downloading it to a local machine, and then provisioning it to all target machines. The README file documents any preinstall requirements including third-party software (e.g., SAS Universal Viewer requires the Microsoft .NET Framework).

In many cases, web downloads include a trimmed-down SAS Deployment Wizard (appearing as a setup.exe, setup.sh, or setup.rexx executable file at the root of the downloaded directory). The README file instructs you to install the software by running the SAS Deployment Wizard from the downloaded directory. You can then use the Record and Playback functionality to create a response file, which can be provisioned along with the downloaded directory in order to install the product on any number of machines in your environment.

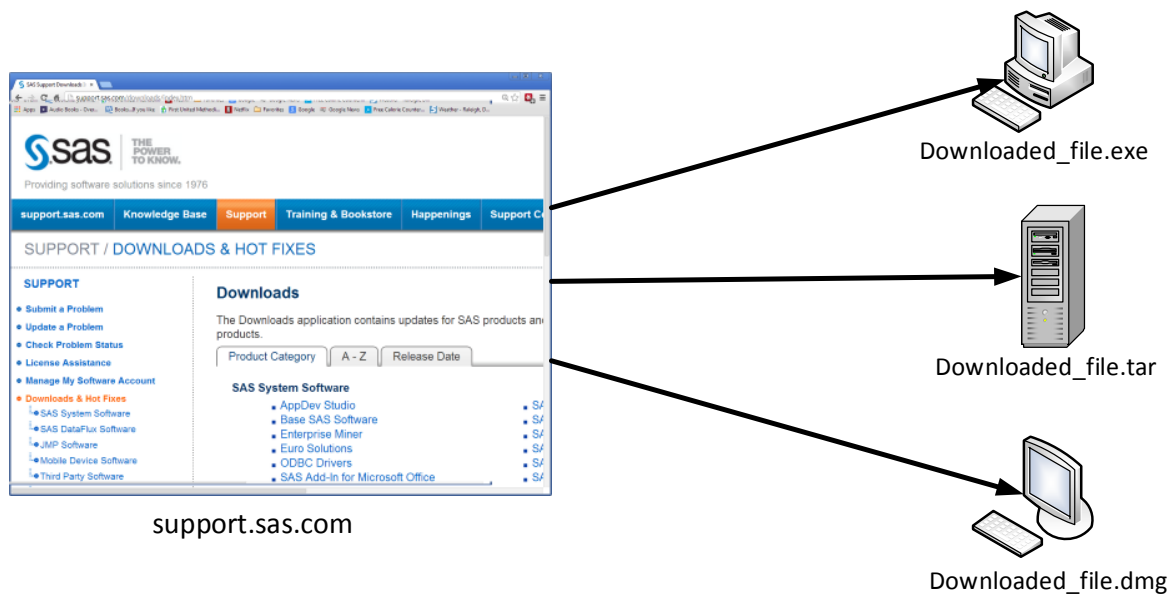


Figure 4. Web Downloads

Advantages:

- Lightweight
- Integrates with existing provisioning software like SCCM and Tivoli
- Includes and installs some, but not all, third-party system requirements (see README file for those not included)
- End users can download and install directly
- Doesn't require target machine access to a SAS Software Depot

Disadvantages:

- Limited product set
- Some system requirements may need to be preinstalled

JAVA WEBSTART

Some SAS software can be run using the Java Web Start (JWS) framework. This technology allows applications to be installed on a web server machine, and subsequently downloaded and launched on-demand within a user's web browser without having to permanently install product files on the local machine. The key advantage is that product updates (new versions, maintenance, or hot fixes) need only be performed in one place - on the server machine. Users receive updates automatically when they launch Java Web Start applications.

The only software permanently installed on a user's machine is the web browser and a version of the Java Runtime Environment (JRE), which supports the application to be launched. Java Web Start technology is built into the JRE. When JWS-enabled applications are launched from a browser, copies of the application's jar files are cached into a user-specific directory, which varies by operating system. Each time the application is launched, the cache is updated with any newer jar files appearing on the web server.

The following SAS products support Java Web Start:

- SAS® Enterprise Miner™
- SAS® Model Manager
- SAS® Forecast Studio
- SAS® Warranty Analysis

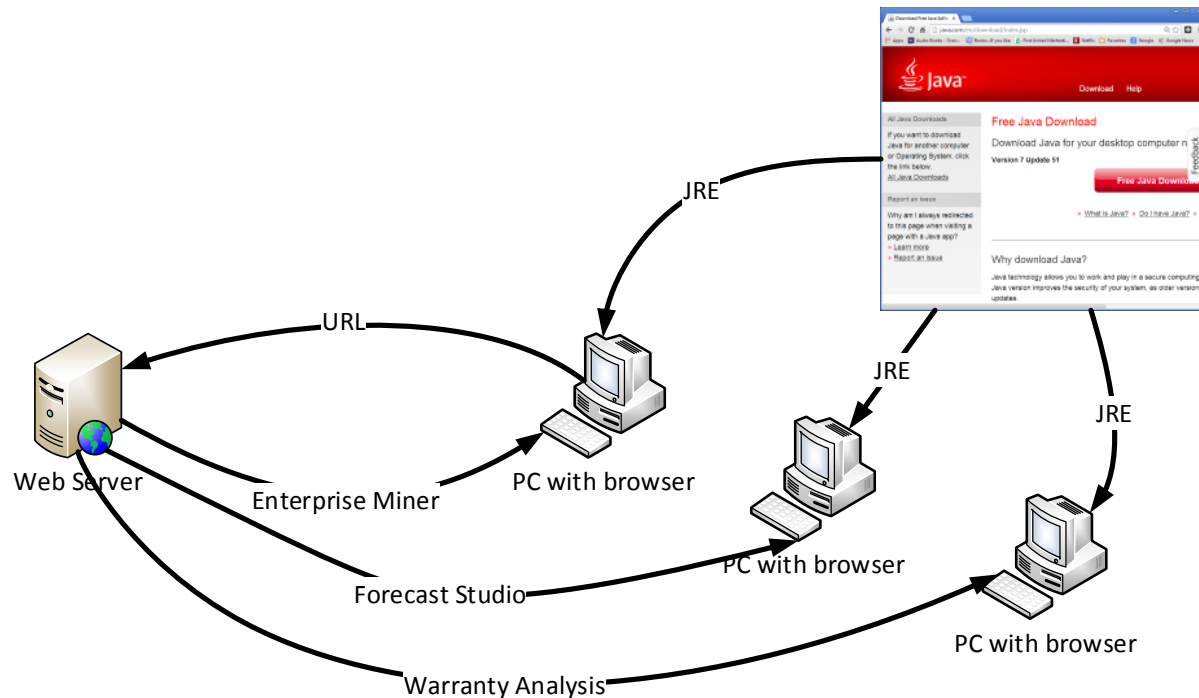


Figure 5. Java Web Start

Advantages:

- Zero permanently installed footprint on target machine
- IT on target machine limited to installing Java Runtime Environment (JRE)
- Doesn't require target machine access to a SAS Software Depot
- Product updates are isolated to a single location and immediately available to target machines

Disadvantages:

- Limited product set
- Requires local Java Runtime Environment compatible with application
- Longer application launch times when product updates haven't been cached

REMOTE DESKTOP SERVICES

Remote desktop services, also known as terminal services, employ machines that can emulate multiple target workstations. Using software like RemoteApp and Remote Desktop Services for Windows or Citrix XenApp, users log in to a remote server, which provides either a full desktop environment or access to specific applications within a shared environment. In either case, the associated SAS client software need only be installed once on the server, or once for each unique virtual machine (VM) maintained on the server.

In a related setup known as Virtual Desktop Infrastructure (VDI), a remote server can host completely separate desktop environments, each with its own dedicated operating system, client software, and user data. While the physical co-location of the remote servers hosting these virtual desktops has certain IT administrative advantages, it calls for separate installation and maintenance of SAS software on each virtual desktop.

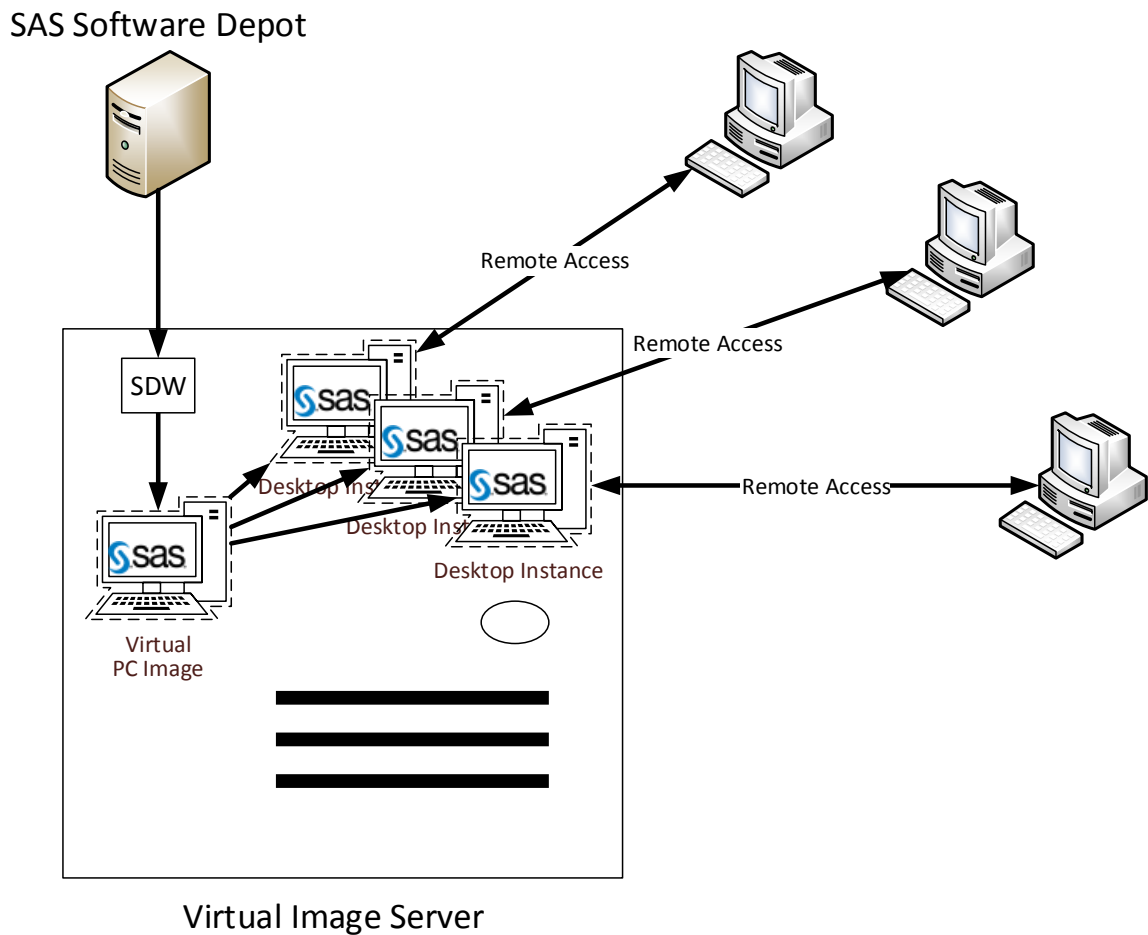


Figure 6. Remote Desktop Services

Advantages:

- Supports all SAS client software
- Zero installed footprint on target machine
- Doesn't require target machine access to a SAS Software Depot
- Product updates are isolated to a single location and immediately available to target machines

Disadvantages:

- Requires network connection to server machine at application runtime (no disconnected mode)
- Requires system resource (CPU, memory, etc.) sharing/balancing on server side
- Is available only on Windows

SAS ONDEMAND FOR ACADEMICS

Colleges and universities can use SAS® OnDemand for Academics to support classes in data management and analytics. SAS Enterprise Guide, SAS Enterprise Miner, SAS® Forecast Server, and SAS® Web Editor are all supported. Users either download and install client software (SAS Enterprise Guide), or access client software through SAS-hosted websites. In all cases, the analytical processing occurs on SAS-hosted servers. Only the GUI runs locally on the target machine.

In the case of SAS Web Editor, provisioning is as easy as notifying users of the appropriate SAS website. For SAS Enterprise Miner and SAS Forecast Server, the user simply needs a Java Runtime Environment in order to support Java WebStart. (See above.) In the case of the locally installed SAS Enterprise Guide, provisioning could involve several of the methods described in this paper.

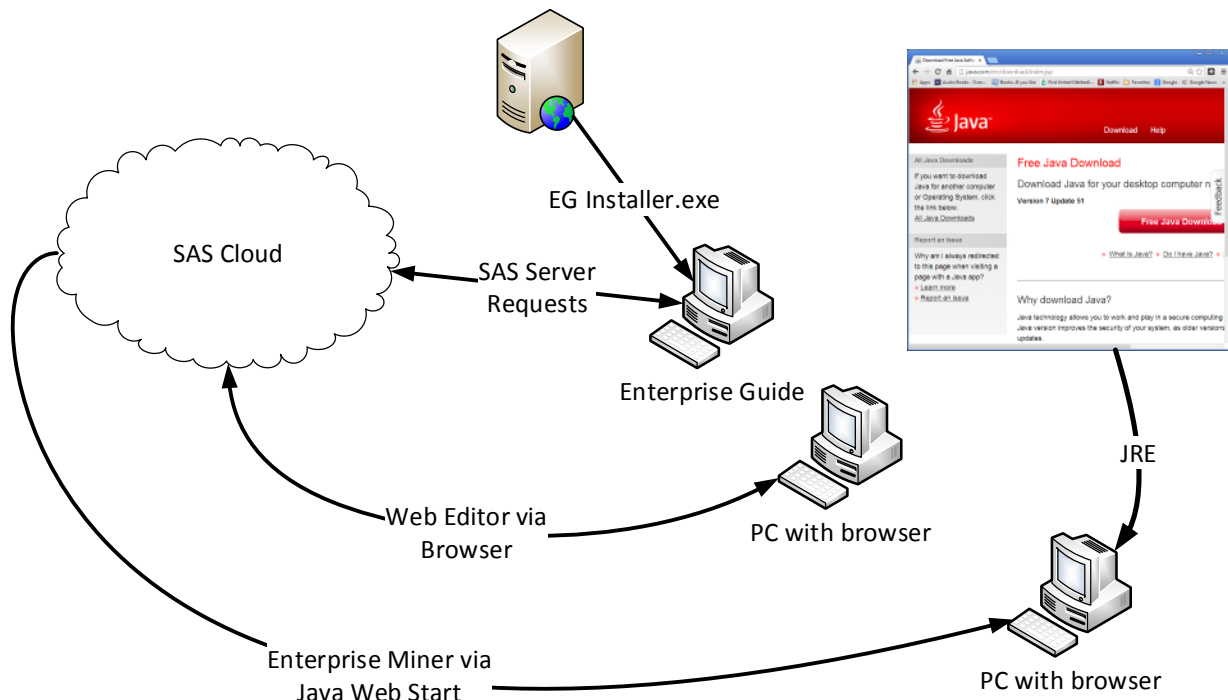


Figure 7. SAS OnDemand for Academics

Advantages:

- No-cost
- SAS hosts server processes in the cloud
- Web-based SAS applications require zero target machine footprint

Disadvantages

- Limited to colleges and universities
- Limited product set (SAS Enterprise Guide, SAS Enterprise Miner, SAS Forecast Server, and SAS Web Editor)
- Requires internet connectivity

APPLICATION VIRTUALIZATION

Application virtualization is a technique involving third-party software like VMWare ThinApp or Microsoft App-V. This software packages up (Microsoft App-V refers to this as “sequencing”) an application into a portable executable that can be distributed and run directly, and in some cases without installing any software on the target machine. The application runs in a virtual “sandbox” that insulates the supporting operating system from change. The sandbox emulates host system resources like files, environment variables, and Windows registry keys. In some cases, virtualized applications can be run on alternate operating systems. For example, you might be able to run a Windows application on a UNIX target system.

Application virtualization allows for easy distribution via web download, network access, or physical media. In the case of ThinApp, it does not require preinstalled software on target machines, and further reduces IT administrative costs by providing single virtualized application images that can be managed in one place and redistributed to multiple target machines. Some applications can be “streamed” to target machines, providing required portions of the application on-demand over the network at runtime.

A thorough description of application virtualization can be found in the 2013 SAS Global Forum paper entitled [“Tips and Techniques for Deploying SAS in an Application Virtualization Environment.”](#)

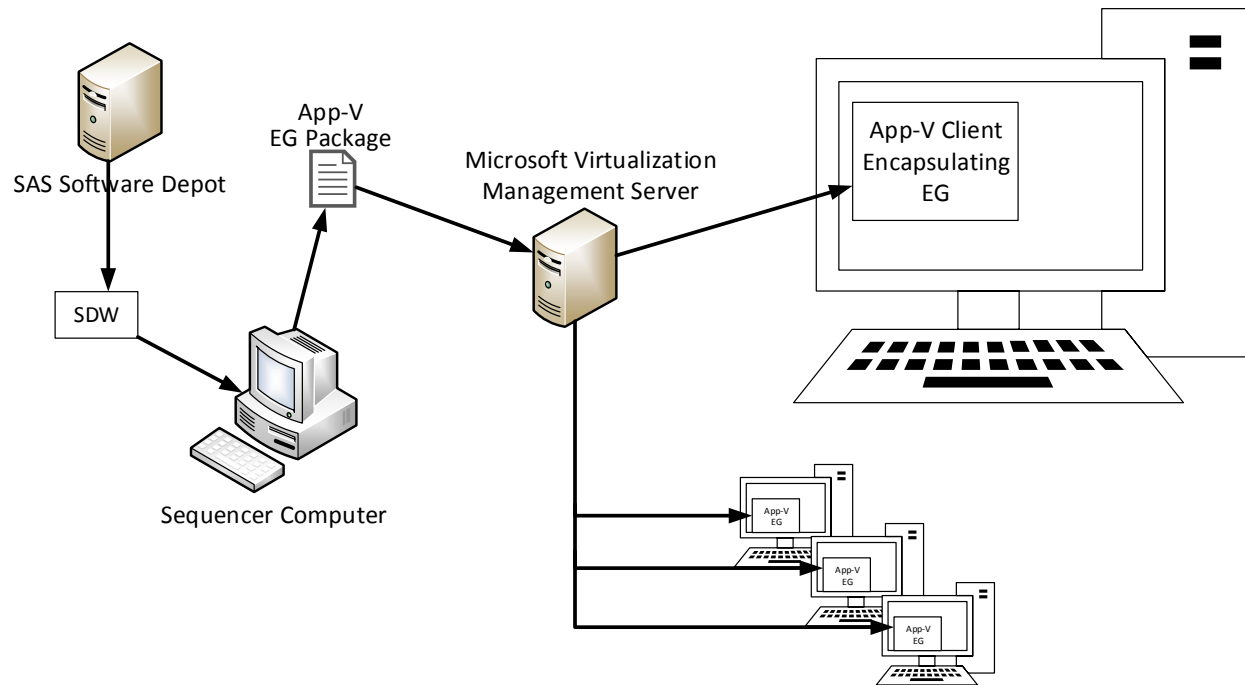


Figure 8. Application Virtualization

Advantages:

- Centralized maintenance of shared virtualized application image
- Minimized impact to target machine by “sandboxing” or encapsulating virtualized applications
- Easy distribution via network, web download, e-mail, or physical media
- Can support virtualized applications for one operating system version running on another
- Supports application streaming

Disadvantages

- Limited product support (typically targeted at desktop software)
- Application updates require redistribution (unless using application streaming)
- Runtime performance hit for the virtualization layer and initial startup
- Still maturing technology that can have trouble sequencing complex applications and supporting all features

SAS SOFTWARE DEPOT SUBSETTING

As SAS software has grown in functionality, it has correspondingly grown in disk footprint. SAS Software Depots, which contain the order or orders you've downloaded from SAS, are particularly large. As an example, a SAS 9.4 Windows order including SAS Enterprise Business Intelligence, SAS Enterprise Guide, and SAS® Visual Analytics, including the latest SAS 9.4 maintenance, takes up approximately 24GB. While this depot is required to deploy the SAS software in your environment, much of its content is appropriate only for machines that will be hosting server software. Client machines require a much smaller footprint.

In some cases, administrators are unable to provide direct network access to these large depots by storing them in central storage locations. In these cases, the depots themselves must be provisioned to other server or client machines from which deployments can be run. However, distributing 24GB depots to multiple machines can be time consuming at best, and logistically impossible at worst. Subsetting these depots can be the answer.

You can trim the size of SAS Software Depots by subsetting them across multiple dimensions – SAS® Foundation products, client products, languages, and operating systems. By default, depots contain a superset of all content for all these dimensions. However, when provisioning software to client machines, it is frequently the case that you require a much smaller subset. Trimming to a single product, language, and operating system can save significant depot size, and make redistribution much more manageable.

As an example, trimming the aforementioned 24GB order to what is required to deploy SAS Add-In for Microsoft Office shaves 75% off the depot size. The remaining depot contains the product content, third-party prerequisites like .NET that support it, and the infrastructure to run the SAS Deployment Wizard itself. Everything else is removed.

While depot subsetting is available when you do your initial order download, it is recommended that you download the full content of your order, and subsequently subset to a second depot using the SAS Deployment Wizard's "Create or Add to a SAS Software Depot" option on the main menu. Since SAS Enterprise Guide and SAS Add-in for Microsoft Office are popular candidates when subsetting depots, the SAS Deployment Wizard provides pre-defined subset recipes for each individually, and both together. Consult the [SAS Deployment Wizard and SAS Deployment Manager 9.4: User's Guide](#) for more information on depot subsetting.

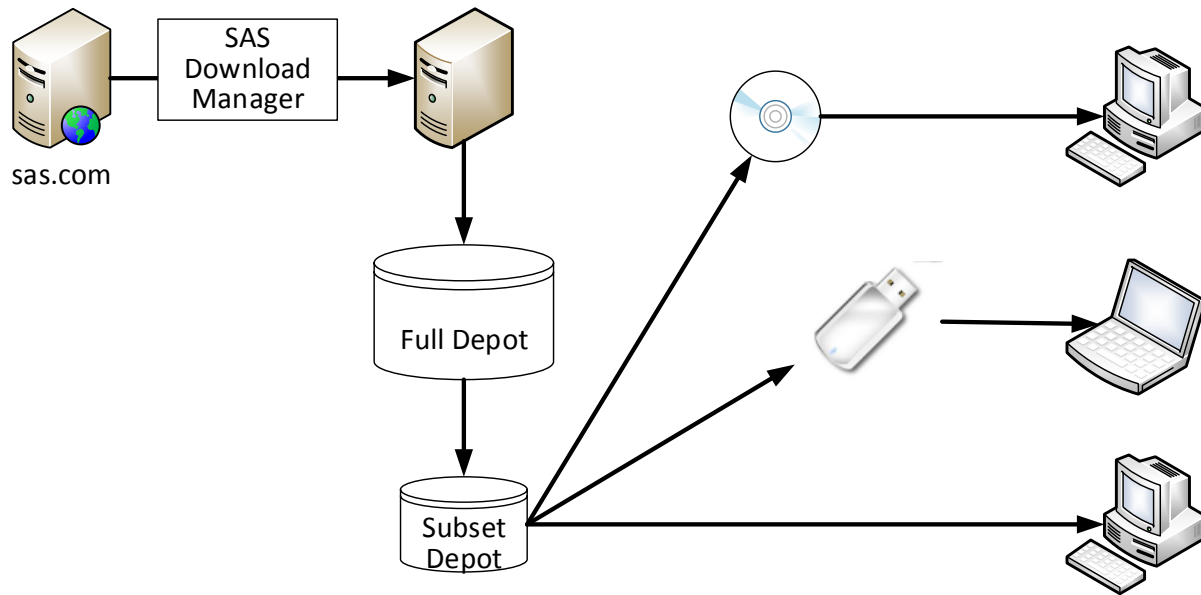


Figure 9. Depot Subsetting

Advantages:

- Supports SAS Deployment Wizard record and playback
- No additional software required
- Subset by one or more dimensions (products, languages, and operating systems)

Disadvantages

- Depots are still gigabytes in size
- Not all products are supported

CONCLUSION

The IT challenge of managing software on hundreds of machines in an environment has motivated the development of a number of cost-saving approaches. Many of these approaches are directly applicable to provisioning SAS software.

Some simplify and automate installation on target machines, like SAS Deployment Wizard record and playback, standalone installers, delta images, and web downloads. Others provide remote or virtual access to client software hosted on server machines, like Java Web Start, Remote Desktop Services, and SAS OnDemand for Academics. Still others involve distributing specially repackaged applications that run in an isolated environment that doesn't impact the target system, as is done with application virtualization.

Finally, SAS Software Depot Subsetting provides a means to reduce depot footprint for easier redistribution when centralized network filespace accessibility is not an option.

There is not a "one size fits all" approach to address this IT challenge. Rather, SAS administrators must assess the available options and pick the one most tailored to the problem at hand. The larger and more complex an environment, the more likely you will be to choose a hybrid approach involving multiple options.

RECOMMENDED READING

"Desktop virtualization." Wikipedia. Available at http://en.wikipedia.org/wiki/Desktop_virtualization. Accessed February 19, 2014.

Hunley, Chuck, Michael King, Max Blake, and Rob Hamm. 2013. "Tips and Techniques for Deploying SAS in an Application Virtualization Environment." *Proceedings of the SAS Global Forum 2013 Conference*. Cary, NC: SAS Institute Inc. Available at <http://support.sas.com/resources/papers/proceedings13/461-2013.pdf>.

Microsoft Corporation. Applications Virtualization. Available at <http://www.microsoft.com/en-us/windows/enterprise/products-and-technologies/virtualization/application/default.aspx>.

Microsoft Corporation. DISM - Deployment Image Servicing and Management Technical Reference. Available at <http://technet.microsoft.com/en-us/library/hh824821.aspx>.

Microsoft Corporation. Imaging. Available at <http://technet.microsoft.com/en-us/library/cc764409.aspx>.

Oracle Corporation. Java Web Start Architecture. Available at <http://www.oracle.com/technetwork/java/javase/architecture-138566.html>.

Oracle Corporation. What is Java Web Start and how is it launched? Available at http://java.com/en/download/faq/java_webstart.xml.

SAS Institute Inc. 2014. *Provisioning Systems and Other Ways to Share the Wealth of SAS*. Cary, NC: SAS Institute Inc. Available at <http://support.sas.com/resources/papers/Provisioning-Systems.pdf>.

SAS Institute Inc. 2013. *SAS Deployment Wizard and SAS Deployment Manager 9.4: User's Guide*. Cary, NC: SAS Institute Inc. Available at <http://support.sas.com/deploywizug94.html>.

SAS Institute Inc. 2006. Revised 2009. SAS Note 17323. "SAS Enterprise Miner Client Deployment Models." Available at <http://support.sas.com/kb/17/323.html>.

SAS Institute Inc. SAS OnDemand for Academics. Available at http://www.sas.com/en_us/industry/higher-education/on-demand-for-academics.html. Accessed February 19, 2014.

SAS Institute Inc. SAS Product Support for Virtualization Environments. Available at <http://support.sas.com/techsup/pcn/virtualization.html>. Accessed February 19, 2014.

ACKNOWLEDGMENTS

Special thanks go to Chuck Hunley and Michael King for their technical review.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author:

Mark Schneider
100 SAS Campus Drive
Cary, NC 27513
SAS Institute Inc.
Mark.Schneider@sas.com
<http://www.sas.com>

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.