

RIS Overview and History

Remote Installation Services (RIS) is a mechanism for deploying remote operating system installations. RIS was introduced in Windows 2000 and essentially provides an infrastructure for automatically deploying operating systems.

RIS Theory

The actual operating system installation component of RIS is nothing new. It is basically either your standard automated install (known as *CD-based image*) or an image download (known as *riprep image*). The integration and potential for complete automation of the entire process is what makes RIS unique.

The basic concept of RIS is that a user can boot up a machine on the network, select a couple options, wait a while, and when it's all said and done have a machine with a fully functioning operating system installed. The user doesn't need any special training or expertise, and the machine doesn't even need to have an initial operating system. The user just boots up with nothing and after a brief period of time is ready to go—like magic.

Of course, this is an oversimplification, but it is a fair approximation of the process. The actual process is initiated when a client boots up on the network and sends a special request. To be able to boot up on the network (especially without an operating system), the client machine needs to have a PXE boot-compatible network card.

Note Most computer systems and network cards are PXE compliant, but if your systems aren't, don't despair. RIS provides a utility called `rbfg.exe` (Remote Boot Floppy Generator) that creates a network boot disk for booting to the network and initiating a RIS installation. The Remote Boot Floppy Generator has drivers for several common network cards.

A PXE boot-compatible network card boots to the network and obtains a DHCP address. The client then queries DNS to find an Active Directory domain controller. The user is prompted for authentication, and the credentials are validated against the domain controller. Active Directory is then queried to determine whether a RIS server responsible for the client exists. If one does, the RIS server is then queried for a list of images that the authenticated user is permitted to install. The list is downloaded to the client machine, the user picks one, and the installation of the chosen operating system begins. When the installation is complete, the client reboots and has a fully installed operating system.

RIS in Windows 2000

As you can probably tell from the process mentioned previously, RIS has several requirements: Active Directory, DNS, and DHCP. Additionally, the RIS server itself acts as a DHCP server and as such must be authorized in Active Directory. Some additional, more subtle requirements also exist for RIS. There must be an NTFS partition to contain the images for

deployment via RIS, and this NTFS partition cannot be the same as the system partition. Configuration of RIS installs a file system groveler that scans the image storage drive to minimize duplicate files to save disk space. Obviously, this could be a bad thing on the system drive.

Implementing RIS is a multistep process. First, the Windows 2000 Remote Installation Service must be installed. As with any service, it is installed from Add/Remove Programs, Add/Remove Windows Components. This just installs the service, however—it doesn't actually do anything. Next, RIS needs to be configured. The Remote Installation Services Setup Wizard, `risetup.exe`, is run to configure the RIS service. It prompts for the location of Windows 2000 Professional source files, copies them to the image storage drive, and creates an initial CD-based image.

Note | Windows 2000 RIS supports only the installation of Windows 2000 Professional. Hot fixes and Service Pack 3 eventually added support for Windows XP and Windows 2000 Server.

Finally, users need to be granted access to download images. Users also need to be granted rights to join the domain or the clients need to be preinstalled.

Limitations of RIS

It's the requirement that users have access to download images that can prove particularly troublesome. Many organizations want to restrict who is allowed to add and remove computers from the domain for security reasons. Because a RIS installation is performed from the client side, the user's account is used to join the domain. This means that the user needs to have the capability to join the domain. If you don't want to grant the user the right to join the domain, an administrator must create the computer accounts ahead of time. The trouble with this is associating the computer account created in the domain with the physical machine that boots up to perform the RIS installation. RIS has a mechanism for making this association, and that is by the network card GUID. Most NIC cards have a GUID associated with them; for those that don't, the GUID is the MAC address with leading zeros. That's all well and good, but how do you know the GUID when you are creating the computer accounts? Ah, there's the rub. You need to physically have the machine available. You can look on the NIC card to obtain the GUID (assuming it's printed there); otherwise, the only way to obtain the GUID is to boot up the machine and press F12, as if booting to RIS. The GUID is displayed on the menu screen.

Another limitation of the RIS in Windows 2000 is that it can be used to deploy only Windows 2000 Professional. As mentioned before, Service Pack 3 finally added the capability to deploy Windows 2000 Server, as well as Windows XP.

When Windows 2000 was released, Microsoft expected organizations to jump on using RIS to deploy it. However, the requirements for RIS—particularly the necessity to have Active Directory—meant most organizations didn't have the necessary infrastructure. Now, with many organizations in the process of or having completed deployment of Windows 2000, RIS is much more feasible for deploying Windows XP and Windows Server 2003.