

Using Virtual Computers to Teach ArcSDE and ArcIMS

Dr. Buren B. DeFee II

Abstract

Teaching the installation and configuration of ArcSDE and ArcIMS can be difficult due to restrictions on administrative access in teaching laboratories. Access restrictions are used to provide a safe, stable computing environment for all users. Removing these restrictions could lead to the disruption of the computing environment or a loss of data.

Virtual computers (VCs) provide a means for students to install and configure ArcSDE and ArcIMS in a safe computing environment without exposing the host computers to undue risk. VCs may be configured to support ArcSDE, ArcIMS and the required software to implement internet-based GIS in a virtual environment. Configured properly, VCs are accessible from the web. This seminar will cover the use of VCs for teaching ArcSDE and ArcIMS in a computing lab which was used to support other courses. Critical configurations for security and access will be covered. Lessons learned will be shared.

Lockdown!

Restricted user privileges are an essential tool to securing computer systems and the networks to which they are connected. However, these same restrictions pose interesting and unique challenges to instructors teaching web-based geographic information systems (GIS). An ideal solution is a dedicated laboratory where the students may be given administrative rights on systems isolated from the other computers on the network. While this may be ideal, a dedicated laboratory takes classroom space and additional computing resources which may not be available. This paper addresses some of the issues and provides solutions to teaching advanced topics to students in a shared computing laboratory.

When I first taught the installation and configuration of ArcSDE and ArcIMS, I utilized a semi-dedicated laboratory. I salvaged a number of outdated computers and placed them into a shared laboratory where their replacements were stationed. Each of the outdated computers were physically reconnected to the requisite equipment - keyboard, mouse, monitor, and network - at the beginning of each class period, displacing its newer counterpart. This process continued until the students installed the operating system and enabled software for controlling the computer remotely over a network connection. Once remotely accessible, the computers were left in a designated area of the laboratory with power and network connections. Further work could then be completed via remote connection.

While this was functional, there were several concerns about the situation. The constant shuffling of computers increased the wear on the systems and placed everything at risk to minor accidents. Since the systems were under constant revision, the network administrator thought it best not to include them in the normal backup process. Students were always in peril of a simple mistake leading to a system failure from which the best recovery was to start over.

Enter the virtual computer

Virtual computing offered a solution to many of these problems. Virtual computing software was used to emulate the functions of a complete computer system within the memory space of a host computer. The data for a virtual computer (VC) is stored in a single file which acts as the VC's hard drive. Once created, the VC operates like a real computer, requiring hardware configuration, installation of an operating system, and installation of software. The monitor of a VC is a window on the host computer (Figure 1).

Virtualization software can support multiple VCs operating simultaneously in a single host. A portion of the host computer's memory is dedicated to each VC, thus the number of active VCs is restricted by the host's hardware. The minimum memory required for each VC depends on the operating system and software installed on the VC.

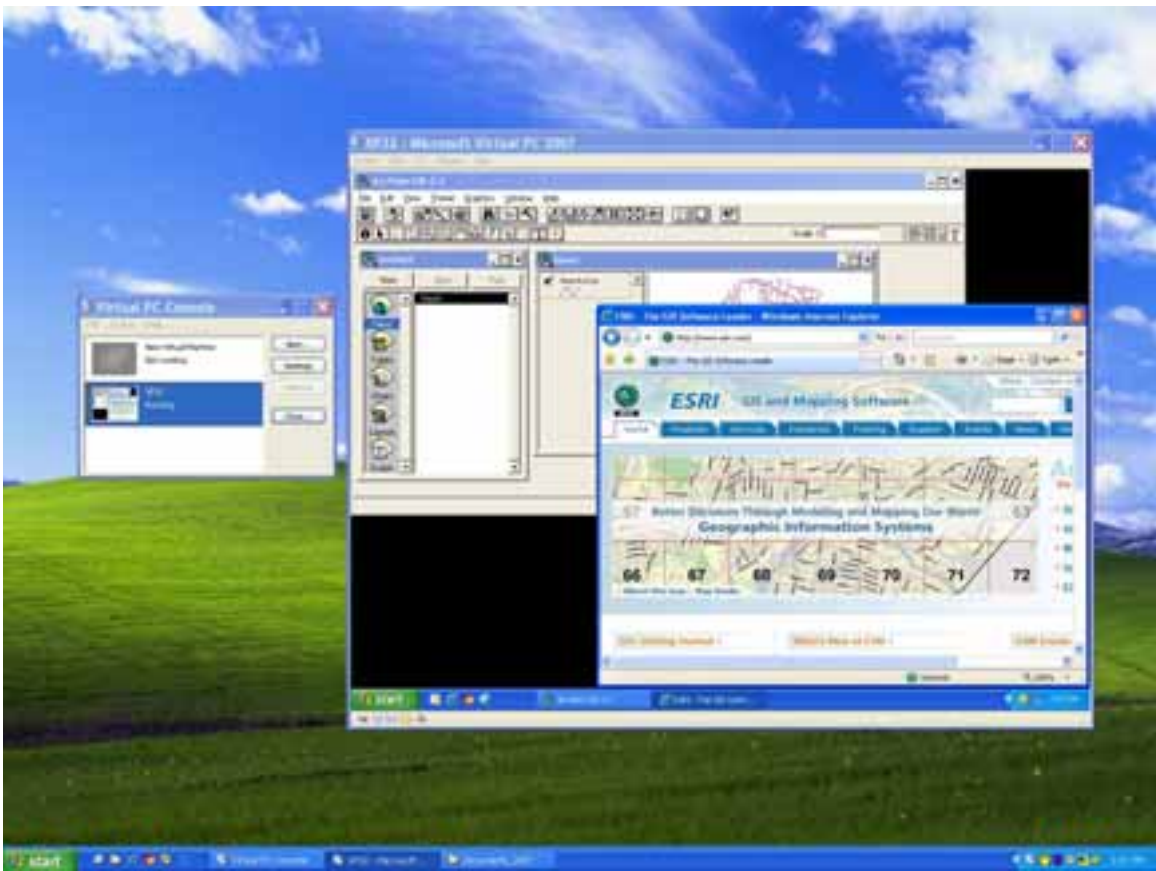


Figure 1. A virtual computer running in a window on a host. This screenshot is a 32-bit XP VC running on a 64-bit XP host.

A VC may communicate to the host, other VCs within the same host, or computers outside of the host, depending on the network configuration. Within a host computer, a virtual network allows for communication. Network address translation (NAT) allows the virtual computer to use the host's network card to connect to the network, but computers outside of the host cannot directly access the VC. This prevents the VC from providing services to the network. If other computers need access to the VC, as in the case of a web-server, a virtual switch on the host routes network traffic intended

for the VC through the host's network card. From the network perspective, the VC appears just as any other "real" computer would. In this manner, the VC may provide any service, including license servers, websites, ArcSDE geodatabases and ArcIMS services.

Keeping it Safe

Virtualization software provides a new level of concern for network administrators. The software enables anyone with user-level access to create a new VC within the network where the user has administrative rights. These new VCs circumvent one of the primary network protections - the campus firewall. It should be noted that it was with great trepidation that the campus Computer Administrator allowed this project to move forward!

To secure the network against rampant virtual computers, access was restricted at three levels. Virtualization software requires an administrator for installation on the host computer. The first layer of protection was installation of the virtualization software into a folder with restricted access. Installed this way, only the administrator, the instructor, and the students enrolled in the course could access and run the virtualization software on the host machine. This prevented the creation and activation of a VC by anyone except specifically authorized users.

The second layer of protection was to store the virtual hard drive(s) in folders to which only authorized users have access. This was easily accomplished by storing VHDs in the user's profile. An alternate location on a secondary hard drive in the host machine was also secured for each student to store their files. Although not required, several students purchased external hard drives to make external backups of their work.

The final layer of protection was within each VC. Once installed, the VC's operating system only allowed authorized users to log on. Students created only two accounts on each VC - an administrator account for installing and configuring software, and a user account for testing configurations. The instructor assigned an administrator password for each student to use. Network administrators were given complete access to the virtual computers so they could shut them down remotely if any behaved erratically.

The access restrictions provided an additional benefit: the VCs could only operate when the student or an administrator was logged into the host computer. This minimized the time that the VC was active on the network, which reduced the exposure to potential outside threats.

Virtual Computing as a Teaching Tool

Many new computers come pre-configured with an operating system and basic software. A new VC is not preconfigured - the virtual hard drive (VHD) hasn't even been formatted! To get an operating VC, the students formatted the hard drive and installed an operating system. The VHDs they created were stored in the aforementioned locations - one for active VCs, the other as a backup.

At each backup, the primary VHD was also "frozen" to prevent further changes. This was accomplished by changing the file properties to read-only. As nearly all operating systems require write access to their hard drives, this freezing process broke the VCs. The virtualization software solved this problem by using a "differencing" disk. A differencing disk is a file stored on the host computer which acts as a portion of the VHD for a given virtual computer. The virtualization software reads both the VHD and the

differencing disk and treats them as a single hard drive. Changes to a VC from the freezing point forward were stored in the differencing disk instead of the VHD. For example, if a VHD was frozen after the operating system was installed, additional changes were recorded in the differencing disk. The virtual computer used the two virtual disks as a single hard drive.

A second advantage of the frozen VHD was the software's ability to use a single VHD to support multiple VCs. Instead of creating a new VHD for each virtual computer, a single VHD with a basic operating system was shared between all virtual computers on a host. Each VC recorded any changes to the basic operating installation into its own differencing disk. This saved disk space and installation time and provided a means for a single, standard VHD to be distributed to each student in the class.

The portability of the VHDs also simplified the grading process. Students turned in completed projects by burning the differencing disk to a DVD or by copying the virtual disks from an external drive into the instructor's computer (Figure 2). The new host computer could reconnect the differencing disks with a standard VHD containing the operating system, thus recreating the VC. Once started, the VC behaved the same as it did on the original host computer.



Figure 2. The virtual hard drives for three virtual computers. The two disks with the "diff" prefix are differencing disks which use the XP32SP2 virtual hard drive for their operating system.

Conclusions

Virtual computing has changed the way I teach my Advanced GIS course. Before virtual computing, each student used an outdated computer which did not meet recommended specifications for any of the software being installed. Students had to physically move and reconnect their computer to external hardware many times during the semester. And students could not utilize the optimum configuration for web-based GIS, one which distributed the software between several computers and offered them the opportunity to better understand network configuration.

Preserving student work was impossible without an external backup system. Failed installations resulted in a reformatted hard drive and the reinstallation of all of the software from the operating system to ArcIMS. Without interruption, this process took up to eight hours to complete. However, in a shared teaching laboratory, it was usually

disrupted by other classes. Students reported repeating this process from three to five times before their final projects were complete.

The versatile nature of virtual computers addressed all of these issues. The outdated hardware initially used for the course was set aside in favor of the new computers in the laboratory. The new hardware efficiently supported two VCs running simultaneously. This allowed the students to configure the software more closely to the suggested installation. They created a VC to house the database and ArcSDE, another to run the web server and ArcIMS. This allowed the students to experience the complications of configuring the virtual systems to communicate over their network and set permissions within each VC. Each VC was also configured to pull data from a computer outside of the host system, providing the students experience with connecting a VC to a "real" network.

Inevitable student mistakes were quickly remedied as backups were readily available. Students were able to freeze a virtual hard drive at any point and copy it to an external drive or burn it to a DVD without complicated backup software. If the installation failed at any step, the student simply copied the previous backup and tried again. Students using VCs reported going through the installation as frequently as previous students, but the time cost was greatly reduced.

The backup process provided portability as well since the VHD could be moved to any system and work could continue when the laboratory was unavailable. One enterprising student did some of his work on his personal laptop. Although this brought up concerns about licensing, the software installed on the VC would only function when the VC had access to the license server which was accessible only through the campus network.

Students turned in their completed work as one or more VHDs. With a few configuration steps I was able to start their VCs on another host. I tested them from a different computer to verify that their complete application worked - from their ArcSDE installation to the websites they built on ArcIMS.

The advantages of using virtual computers include increased security, easier backups, less reinstallation time, improved student understanding of networks, less disk space consumed in multiple configurations, and easier grading for the instructor. Virtual computers have eliminated the need for dedicated hardware or a dedicated laboratory for teaching advanced topics in GIS. Its ease of use has opened the door to exploring alternate configurations for web applications in future classes.

Acknowledgements

I would like to thank Paul Freeman, IT Manager for the School of Forest Resources, University of Arkansas at Monticello, for his invaluable help in implementing this project.

Author Information

Buren B. DeFee II, Ph.D.
Assistant Professor of Spatial Information Systems
School of Forest Resources
University of Arkansas at Monticello
P.O. Box 3468
Monticello, AR 71656
(870) 460-1748
defee@uamont.edu
buck@defee.com