The Virtue of Virtual: Using Virtualization Technology to Deploy Training Environments

by Robert D. Brown, RDB Consulting, Inc.

Introduction

We are now solidly into the era of computer system virtualization. Virtualization allows you to create a total computer environment consisting of operating systems, applications, databases, and servers, and then to run this environment on any host PC without affecting any of the existing applications or settings on the host PC.

Although there is a bit of a learning curve involved with setting up virtual training environments, and then deploying these environments, the beauty of using a virtual approach is that you only have to do it once, instead of having to configure each machine in your classroom.

Problem Overview: Conventional Training Environments

Let me first give an overview of the problem and then explain how virtualization provides a solution. It used to be that applications were relatively straightforward and not too large in size. If you wanted to train someone on a particular application, all you had to do was install a copy of the application onto each training PC or laptop and you were all set. Each student had a simple set of data values that they plugged in as they did exercises, so that when class was over, you just uninstalled the application and reinstalled a fresh version for the next class and you were all set.

This approach was fine with stand-alone applications. However, with enterprise applications using directory servers and databases, setting up training PCs is considerably more complex. These enterprise applications take much more time and expertise to install and configure on each PC, and the application on a student's PC is just one component of a system that consists of networked servers and databases.

Typically, the solution to setting up a training class has been to set up a single training environment on a dedicated training server. The required databases directory server and web servers are all installed and configured on the dedicated server, and all students in the class access this single training server. Depending on the application and the training classroom setup, students access the training server through the internet, a LAN, or a VPN.

With many enterprise applications, there is an application installed on the student PC (front end) that accesses the servers and databases (backend). Increasingly though, the trend is to build an internet-based front-end, so that no application other than a standard Internet browser is needed to use the application. In this case, students use the internet and access a single instance of the application, logging on with a unique user name.

As can be seen from the above descriptions of conventional training environments, there is a great deal of technical complexity involved in setting up a training classroom to train students on networked enterprise applications.

The complexity goes beyond merely establishing secure and reliable network access, however. While this aspect is difficult enough, the complexity is further compounded by the nature of the training itself. If all students are completing the same lessons, and they are all accessing the same database, they cannot all create a user named, for instance "John Doe." A typical work around for this problem is to assign each user in the class a number, so that all values the student enters have this suffix.

For example, when user 01 enters the value *John Doe*, he enters it as *John Doe01*, user 2 enters *John Doe02*, and so forth. While this approach is workable, the probability of human error is high. Whenever students forget to enter the suffix or enter a wrong suffix value, they inadvertently make changes to one of their classmate's inputs.

This approach also raises significant maintenance issues as the training server databases must all be purged once a class is completed. As anyone who has worked with networked enterprise applications will tell you, this task is never easy, and it is easy to accidentally leave some old data in the system. If the old values are not completely purged from the system, this causes problems for the next set of students.

Another much more expensive approach that some companies take is to have a bank of servers dedicated to training so that there is a one-to-one relationship between each student and server. This solution is more effective than the many-students-to-one-server approach discussed above, but maintenance after each class conduct is required and this maintenance is multiplied by the number of servers being used. This is also a very expensive approach in that many servers and many versions of the supporting software (i.e., Web application servers, databases, directory servers, etc.) are used, and software licenses must be purchased for each. The probability of one of the servers crashing or becoming disconnected is increased.

Until recently, there was no real way out of this dilemma. As applications have become more complex, the difficulties and risks associated with deploying training have increased exponentially.

Virtualization Solution

However, in the last several years, several virtualization products have entered the market that make training deployment much easier. The two market leaders as of writing are EMC's VMware® (www.vmware.com) and Microsoft Virtual PC® (www.microsoft.com/windows/virtualpc/default.mspx). Note that neither product is compatible with a Mac OS, so for the present the remainder of this discussion pertains to computers running on a Microsoft Windows (XP or Vista) platform. (VMware® supports the Linux operating system in addition to Windows, but this is outside the scope of this white paper.)

Here is how virtual training environments can be used for delivering training. Using one of the virtualization products just described, you build and configure your training environment on a

single PC or server. This training environment typically includes the application itself, and any Web application servers, databases, directory servers, API connectors, or other necessary components. The virtualization product saves all of these applications and associated configuration and data as a single large "image" or file that can be saved and copied. You can copy this file onto a CD (if small enough) or more likely onto a DVD, or place it onto an FTP file server, where it can be downloaded.

To deploy the virtual training environment, you must first install a small (less than 10MB), client version of the virtualization software on each student's PC. Each student then downloads or copies the image of the virtual training environment to their PC's hard drive. Once copied, they can open the file using the client version of the virtualization software. This opens a window on their PC that allows access to the complete training environment with the application servers and databases all connected and functioning correctly. It is as if they have remote access to external databases and servers, except that the entire training environment is encapsulated on their PC.

Advantages

The advantages of this approach are many. Every student can enter the same values without worrying about corrupting other students' information. There are no concerns about connectivity, network access security, or response times. The training is completely portable. It can be conducted anywhere students have laptops or PCs, as there is no need for network connections.

There are almost no maintenance issues with this approach. Once the class is over, the virtual training environment file is deleted from each student PC.

Another advantage is that if a student makes a major error with the application during the training and cannot recover, the instructor can simply delete their current virtual training environment and reinstall it, so the student can start again with a clean slate.

Similarly, instructors can keep master versions of the training environment with various degrees of the exercises completed. For example, if a course is two days long, an instructor will usually have a master with all of the first day's exercises completed. This way, if a student falls behind, the instructor can simply install the master with the first day's exercises completed so that the student can begin the second day of class along with everyone else.

Limitations

Like any relatively new technology, virtualization is not perfect. There are some limitations—the first being size of the file. This really is a limitation regarding current media storage capacity and not of virtualization itself, but it has a practical impact. Maximum practical capacity for a DVD is about 4 GB, so if you are planning to deploy via DVDs, keep this size limitation in mind when building your training environment.

Download Times from FTP Site

As stated above, there is no inherent size limit to your virtual training environment, so theoretically you could have a virtual training environment with a size of 100 GB, but again, you must keep in mind the practical consideration of how long will it take for the student to download the file to his or her PC.

Student Storage Capacity

Also keep in mind that each student's PC must have sufficient space on its hard drive to hold the large virtualization file. If, as is increasingly common, students are bringing their laptops to class, you must notify them in advance of the space requirements on their hard drive. A good rule of thumb is to require that the student's PC must have at least 15GB of free space on its hard drive. This is based on the assumption of a 4GB virtual training environment, leaving an additional 10GB or so free so that system performance is not degraded due to an overfull had drive.

Memory Capacity

Both the leading virtualization products mentioned in this white paper can run on as little as 300MB of RAM, but for optimal operation at least 1GB of RAM is preferred and 2 GB of RAM allows for almost seamless operation, letting the student toggle between the virtual training environment window and their PC with minimal disruption.

Processor and Operating Systems

Both VMware® and Microsoft Virtual PC® are very specific about their requirements for processors and operating systems. You can check the documentation for each product for details. Compatibility is improving with each new release of both products, but a good rule of thumb is that each student PC should have a 486 or above Pentium processor, and a Microsoft XP or Vista operating system. If you have using PCs with processors or operating systems other than those just mentioned, carefully research the virtualization product's documentation regarding compatibility. Both products are very unforgiving in this regard. If you try to run either VMware® or Microsoft Virtual PC® on a PC with an incompatible processor or operating system, you risk a major system-level crash that could cause permanent damage to the PC.

Student PC Disk Drives

If you plan to copy the virtual training environment to each student PC from a CD or DVD, each student PC must have a CD/DVD drive. Although most PCs come thus equipped, a sizeable minority do not. In this case, a good strategy is to provide your instructors with a small portable external had drive so that they can copy the virtual training environment file to those student PCs without a CD/DVD drive.

Conclusion

This ability to virtualize computer environments is transforming the IT industry, and has the potential to transform IT training as well. Virtualization eliminates many of the training delivery

problems training professionals have been plagued with as applications have grown larger and more complex. Costly, time-consuming training setup and maintenance issues are eliminated, and students are not subject to slow response times or network connectivity problems. As a result, your training classes will flow smoother, students will have less distractions and interruptions, and costs associated with training delivery will be dramatically reduced.

For more information about virtualized training environments and how they can dramatically reduce your training costs while increasing reliability and security, please contact Rob Brown at RDB Consulting, Inc. by phone:1-214-684-1263 or e-mail: rob@rdbconsulting.biz

About Robert D. Brown

Robert founded RDB Consulting, Inc. in 2001 to focus on developing product training and communications solutions for technology-driven companies. Several years as a journalist honed his reporting and writing skills and led to a successful career in training and communications for large high-tech and engineering companies. Prior to founding RDB Consulting, Robert was an Accenture Consultant in their Change Management practice. He has also held positions with Dallas Area Rapid Transit (DART), the U.S. Department of Energy Superconducting Super Collider Laboratory, and Texas Instruments.

About RDB Consulting, Inc.

RDB Consulting, Inc. has been in business since 2001, and has established a proven record of success in designing, developing, and implementing award-winning training, documentation, and communications for such clients as:

- *Cisco Systems*
- *Hewlett-Packard*
- *Voyence*
- *Fuego*
- *TruLogica*