Teaching Tip

Using Virtual Servers to Teach the Implementation of Enterprise-level DBMSs: A Teaching Note

William P. Wagner, Ph.D.
Accountancy and Information Systems
Villanova School of Business
Villanova University
William.wagner@villanova.edu

Vik Pant
Oracle Corporation
Vik.pant@oracle.com

ABSTRACT

One of the areas where demand has remained strong for MIS students is in the area of database management. Since the early days, this topic has been a mainstay in the MIS curriculum. Students of database management today typically learn about relational databases, SQL, normalization, and how to design and implement various kinds of database applications. Choosing the best database package on which to base the hands-on activities for the class is one of the biggest difficulties of teaching it. Many schools end up using the lowest common denominator which is often a PC-based package such as MS Access. While this package is certainly adequate, basing the whole course on using this one package may be a disservice to the students. This is because in the real business world, the market is dominated by the larger, enterprise-level DBMSs such as Oracle, IBM DB2, and Microsoft SQL Server. However, getting the students access to one of these packages poses a whole series of other problems involving getting academic alliances and maintaining it on a server. This article describes a novel use of a server virtualization tool that is commonly used in the business world to allow students to gain experience in using several of the most popular enterprise-level DBMSs. Students benefit because they get more real-world experience out of the class and gain confidence that their skills will more readily transfer to jobs. Faculty also benefit by having more flexibility in how they can present the material.

Keywords: DBMS, IS curriculum, server virtualization, database education, virtual server

1. INTRODUCTION

One of the main functions of IS faculty is to carefully choose topics for the IS curriculum that will give students the most powerful and complete set of skills and knowledge within the constraints of the available resources. Today's workplace is certainly dominated by enterprise-level DBMSs from vendors such as Oracle, IBM, and Microsoft. Having skill and experience with one of these enterprise-class DBMSs can greatly enhance a student's chances of a successful entry into the world of IT careers (Urban and Dietrich, 1997). In fact, the US Bureau of Labor statistics for 2008-9, indicate that database administration is one of the fastest growing industries in the entire US economy (www.bls.gov/oco/ocos042.htm#outlook). It further states that this should be one of the fastest growing occupations through 2016.

The problem then for IS educators is how to best prepare students to enter the overall IT job market with database skills. For many schools, purchasing a dedicated server and the required licenses is beyond their resources. Or they may not have the in-house skills necessary to maintain a server. This paper proposes one alternative solution to this problem of educational resources. In this case, students obtained experience in installing and using a total of four different enterprise-class DBMSs; including Oracle, DB2, SQL Server, and MaxDB. This was made possible by the innovative use of server virtualization software from VMware. Thus, students received valuable hands-on experience with powerful DBMS packages, as well as with a popular server virtualization tool. Reports from alumni working in the field indicate that this is also a highly-valued skill in the corporate world today.

2. WHAT IS SERVER VIRTUALIZATION?

Server virtualization refers to the process of partitioning the physical resources of a computer server into sub-sets

and apportioning them to logical, or virtual, systems. This approach is achieved by the abstracting the underlying hardware via a control layer in between the physical server and the virtual servers. In the software world, the use of multiple abstraction layers helps to hide the complexities of the previous layer in order to make the particular operating system device independent. Similar abstraction layers can be seen in communication protocols such as the popular OSI-7 layer network model. In the context of server virtualization, the control layer is referred to as the virtualization engine or hypervisor. The physical server on which the virtualization engine runs is referred to as the host system. Each of the virtual systems that in-turn run on the virtualization engine are called guest systems. The term "virtual" is provided to servers used in this manner because, unlike a physical server with full control of all of its physical resources, these logical systems make use of the physical resources present in the underlying host system.

Figure 1 exhibits a sample schematic of a physical server that functions as a host server to a number of virtual systems via a virtualization engine. Since the physical resources in the host server are finite, this in turn limits the resources available to the guest servers on that host server. The virtualization engine can only supply resources to the virtual systems it supports within the limitations of the operating system of the host server. Typically, host servers can be configured to optimize the allocation of these resources used by the virtual systems in order to improve the overall system performance.

There is a running debate on the relative merits of two different approaches to server virtualization. In the classic or Type 1 approach, the virtualization engine, or hypervisor, runs directly on top of the hardware. Examples of Type 1 server virtualization include IBM's z/VM (originally CP/CMS) and the more current Citrix software. VMware is an example of the popular Type 2 server virtualization software, where the engine runs on top of the host operating system, and the guest systems run on level three. Some view the Type 1 approach as more secure while the Type 2 approach gives administrators more flexibility.

3. PRESENTING ENTERPRISE-LEVEL RDBMSS USING VIRTUAL SERVERS

In the past few years, a number of IT educators have begun to explore how to enhance their teaching using virtual servers (Albee et al., 2007; Bullers, Burd, and Seazzu, 2006; Gaspar et al., 2008; Pons, 2003; Stackpole et al., 2008; Stefan and Abu-Mulaweh, 2009;, Vollrath and Jenkins, 2004). Most of these attempts were in the context of teaching students the various concepts of network administration (Albee et al., 2007; Bullers, Burd, and Seazzu, 2006;,Gaspar et al., 2008; Stackpole et al., 2008; Stefan and Abu-Mulaweh, 2009; Vollrath and Jenkins, 2004), though some have extended the reach of server virtualization into cyberwarfare (Stefan and Abu-Mulaweh, 2009), and e-business in general (Sandvig, 2007). Only one researcher has actually proposed using it for teaching database technology (Bullers, Burd, and Seazzu, 2006). But all seem to agree that teaching with and about server virtualization offers a variety of student and faculty benefits. This paper is the first to propose how the two topics can be combined in an innovative way.

In a typical Database Management class at Villanova University, there will be somewhere between 20-25 students. This three credit course meets twice a week and is required of all MIS majors at Villanova. One unique feature of Villanova's business school is that all students are issued new Dell laptops every two years. These laptops come pre-loaded with an image that includes Microsoft Office Professional with MS Access. Students learn the basic principles of relational databases and how these are implemented in MS Access in the first half of the semester. Then, in order to gain more experience with enterprise-class DBMSs, they learn how to install their own copies of DB2, Oracle, SQL Server, and MaxDB.

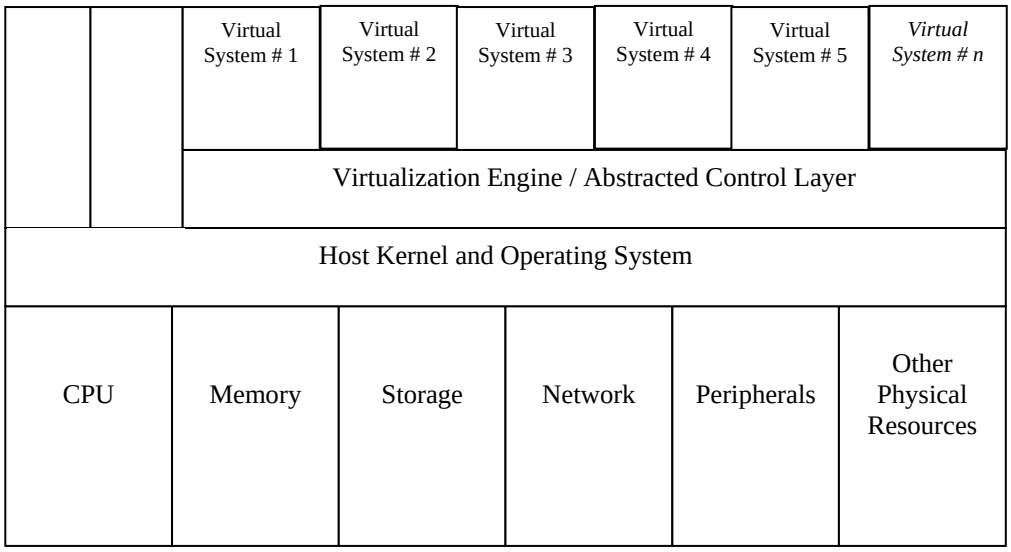

Figure 1: Virtualization schematic

In order to do this, a basic dual processor server was obtained along with a license to a popular server virtualization tool, VMware. VMware was one of the originators of Type 2 server virtualization software and is currently the market leader. Now, all the major software vendors such as Microsoft, IBM, and Oracle, also have competing server virtualization products. In general, the server performed adequately, though it was found that if more than 4 or 5 installations were being attempted at once, the system response was so slow as to be unworkable. This is why it was thought to be best to structure the exercise as a team project with no more than 5 teams.

In Villanova's experiment with server virtualization, two students volunteered to maintain the server and to create the images needed for this project. An alumnus of the MIS program at Villanova who had industry experience in using VMware trained the first two students. They then worked together to create four separate images; one for each of the enterprise-class database packages presented to the class. These images included the appropriate operating system, 90-day registered versions of the DBMS, along with all the hot and cold patches necessary to get the DBMS up and running quickly. The basic VMware console is shown in figure 2. It shows three base images running Windows 2003 Server, XP, and Linux respectively.

At the beginning of the first class, a brief overview of the various enterprise-class DBMSs was presented and the installation of one of the packages was demonstrated by one of the team leaders. Students were then divided up into separate groups of four or five students and one of them was also designated as the team lead. These teams were then assigned to install one of the four DBMS packages: DB2, Oracle, SQL Server, or MaxDB (a free DBMS used by SAP). Then, each of the team leaders connected to the class server and logged into the appropriate image that had been created earlier. By following basic instructions, each group managed to install and run a complete DBMS on this server by the end of an hour and fifteen minute class. In following weeks, each group rotated to a different DBMS and went through the installation process for that database package.

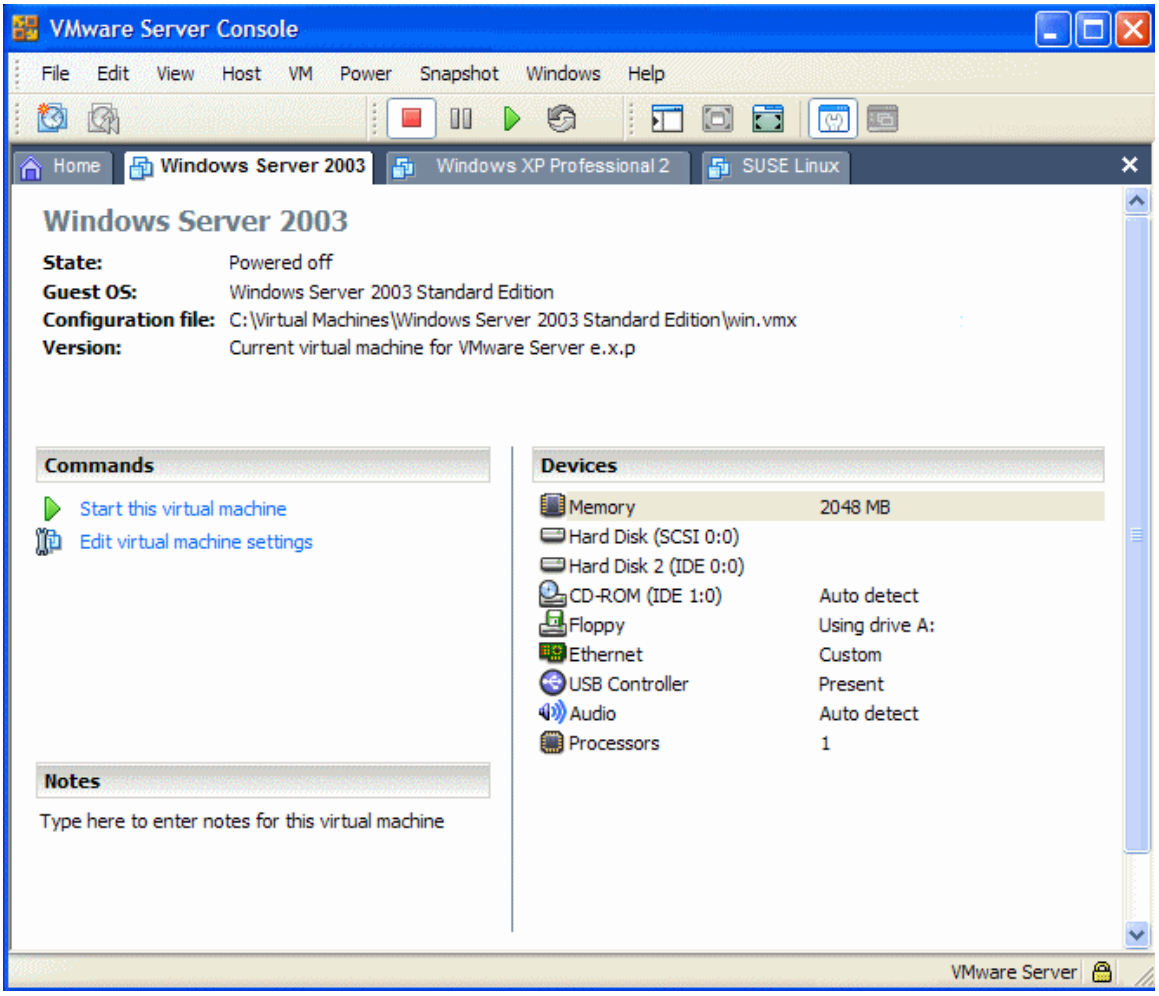

Figure 2: VMWare console with three base images

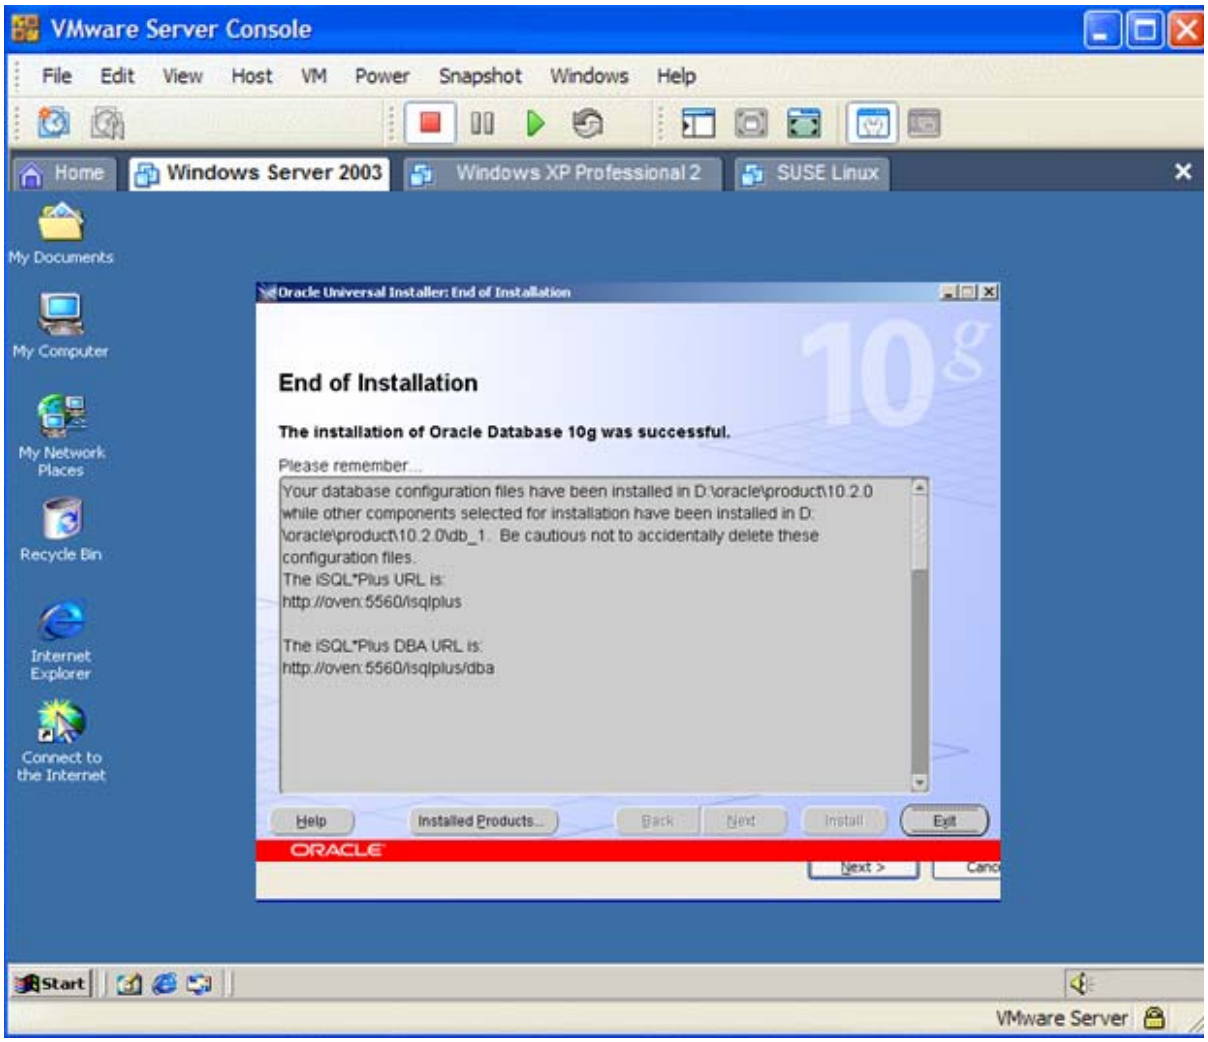

Figure 3: Installation of Oracle running on VMWare image

Figure 3 shows a screen shot of a shareware version of the Oracle installation software that was downloaded from Oracle's public website.

This approach meant that all the students had a wide exposure to a variety of enterprise-class database packages and they gained confidence in using them. After four weeks, each group had installed each DBMS once and each student had also been a team lead for the installation once. Other members of the group were required to document the installation process and research unique features of each package. Students later went on to install their own copies of the various DBMSs on their own machines as they needed to refresh their knowledge about them for upcoming interviews and other class projects.

4. BENEFITS OF USING VIRTUAL SERVERS IN THE CLASSROOM

Using virtual servers to help teach students about how to implement an enterprise-class RDBMS has a number of benefits. Not only do students expand their knowledge about the similarities and differences among the top database packages, but working with virtual servers helps to reinforce other important concepts. These include the ideas of reuse and base-lining, adaptive resource allocation, rollbacks, and general network database administration concepts. These benefits are discussed in more detail in the following sections.

4.1 Ease of preparation and Reusability

Virtual servers make the task of initializing or base-lining systems relatively easy ahead of a class meetings. Initializing refers to the process of preparing the server with basic tools and utilities and implementing an operating system on it. Base-lining refers to installing the components that are pre-requisites to the lecture that will be presented in the classroom. We found that it is difficult to demonstrate and teach complicated installation procedures of enterprise-class database systems (or any other type of enterprise-class software) in a classroom environment because the burden of initializing and base-lining a system is too time-consuming. For example, an operating system needs to be installed and configured on each server on which the software will be installed. Then each of the other prerequisites for that software also needs

to be installed on the server. This usually includes components such as the Java Software Development Kit, ODBC connectors, JDBC drivers, custom adapters and various interfaces relevant to the software being installed. These initialization and base-lining processes often take a long time and so regardless of whether these are performed in front of the audience or ahead of the presentation they are not often directly relevant to the installation of the software that needs to be taught.

Using server virtualization software in the classroom provides the capability for instructors to create mirror images of guest servers and enable them simply by copying them to a new location on the file system and then performing some basic customizing. This is a major time-saver and means that a guest server can be readied to the point that just the installation of the software (which is to be demonstrated) is remaining. At that point it can be copied as many times as needed to demonstrate the installation process over and over again.

4.2 Dynamic resource allocation – Contextual and adaptive

The virtualization engine resides between the physical resources of the host server and each of the guest virtual servers. By design, the virtualization engine controls the resources that are allocated to each of the guest servers from the resources that are provisioned to it from the host system. By being able to control how the resources that are available to the various virtual servers, the virtualization engine can dynamically adapt to the needs of the various software packages running on the guest systems. For resource-intensive applications it can monitor the context and usage and add or subtract memory and processing power as needed. The importance of this concept for students who are learning about database technology has been noted by at least one other researcher (Urban and Dietrich, 1997).

4.3 Rollback and reversion

Many of the leading server virtualization products, including VMware and Microsoft's Hyper-V, come with the capability to capture snapshots or copies of the state of a server at a point in time and save this to the disk for later retrieval. In the corporate world, this is a very valuable feature which is rarely mentioned in classrooms, but is one of which students should be aware. Students learn that if a person makes a mistake in one of the installation steps and the installation fails then it is possible to revert back to a snapshot of a prior, safe point in the installation sequence. Although this is possible it is not easy in a standard operating system running on a typical server.

4.4 Ease of administration

Virtual servers typically turn out to be much easier to administer and manage than having multiple, independent servers. The time and costs associated with managing the physical aspects of the multi-server environment are typically much higher than compared to using server virtualization. Even though the virtual servers are provisioned physical resources via the virtualization engine, there is in fact, only one physical device to manage and maintain. This fact has led to a huge growth and proliferation of server virtualization in industry today.

Server virtualization also simplifies server administration by making it easier to roll out system patches, updates, and hot fixes to the entire enterprise. Rather than fixing multiple servers individually, patches and updates can be distributed across multiple virtual using the base images on each of the guest systems. These changes and updates can be cascaded to all the virtual servers as needed in a quick and efficient manner.

4.5 Monitoring and control

Often in classrooms it becomes difficult to monitor what the students are doing on the systems on which they are working. Even when students are logged onto the servers via remote desktop or other screen-sharing sessions, it is difficult for the instructor to keep watch over each and every student. Most virtualization software products have built-in tools for network monitoring and control. Within VMware, administrators have the ability to "scroll" through the desktops that are running on that host server. This is helpful for the instructor because not only can he or she view each student's desktop, but they can also control their interactions with the host system.

5. CONCLUSION

This paper describes an innovative way of teaching students how to install and use a variety of enterprise-class database management systems. It solves the problem of how instructors can get students working on high-powered database packages while simultaneously educating them about how to implement and use popular virtual server technology. This has not been proposed with specific details by other researchers and IS educators to date. Students derive a variety of benefits from this experience and also gain confidence in working with real-world systems that are omnipresent in the world of information technology today.

6. REFERENCES

Albee, P. , Campbell, L. , Murray, M. , Tongen, C., Wolfe, J. (2007), "A student-managed networking laboratory," *Proceedings of the 8th ACM SIGITE conference on Information technology education*, October 18-20, Destin, Florida, USA.

Bullers, W., Burd, S., Seazzu, A. (2006), "Virtual machines - an idea whose time has returned: application to network, security, and database courses," ACM SIGCSE Bulletin, v.38 n.1.

Gaspar, A., Langevin, S., Armitage, W., Rideout, M. (2008), "March of the (virtual) machines: past, present, and future milestones in the adoption of virtualization in computing education," Journal of Computing Sciences in Colleges, v.23 n.5, p.123-132.

Pons, A. (2003), "Database Tuning and its Role in Information Technology Education." Journal of Information Systems Education,, Vol. 14, Issue 4.

Sandvig, J. (2007), "Selection of Server-Side Technologies for an E-Business Curriculum." Journal of Information Systems Education,, Vol. 18, Issue. 2.

Stackpole, B., Koppe, J., Haskell, T., Guay, L., Pan, Y. (2008), "Decentralized virtualization in systems administration education," *Proceedings of the 9th ACM SIGITE conference on Information technology education*, October 16-18, Cincinnati, OH, USA.

Steffan, G., and Abu-Mulaweh, H. (2009), "Teaching Local Area Networking in a Secure Virtual Environment." Computer Applications in Engineering Education,, Wiley online.

Urban, S., and Dietrich, S. (1997), "Integrating the practical use of a database product into a theoretical curriculum." ACM SIGCSE Bulletin,), Vol. 29, Issue 1.

Vollrath, A., and Jenkins, S. (2004), "Using Virtual Machines for Teaching System Administration." Journal of Computing Sciences in Colleges,Vol. 20, Issue 2.

AUTHOR BIOGRAPHIES

William P. Wagner has been teaching, consulting, and conducting research in Information Systems since 1988. He has been on the faculty of Villanova University since 1991 and has developed many innovative new courses and been a leader in the use of new teaching technologies. His research includes topics in the areas of Artificial Intelligence, Enterprise Resource Planning (ERP) systems and E-Commerce and has appeared in internationally known journals such as *Expert Systems*, *Journal of Computer Information Systems*, *Internet Research*, and *Knowledge-Based Systems.* He recently co-authored the first university-level text book on eXtensible Markup Language (XML) entitled *Introduction to Applied XML Technologies for Business* (Prentice Hall, 2003), and the first text on Customer Relationship Management (CRM) entitled, *Customer Relationship Management (CRM): A People, Process, and Technology Approach (Thomson Publishing, 2007).*

Vik Pant is professionally engaged with Oracle's sales organization. His current focus is on Oracle's Fusion Platform and its Services Oriented Architecture. Vik's experience encompasses business development and sales execution centered on emerging technologies. He is an avid researcher and his academic research has been published in numerous refereed scholarly journals. These include the *Journal of Strategic E-Commerce*, the *Journal of Electronic Commerce in Organizations*, and the *International Journal of Enterprise Information Systems*. He also serves on the review boards of the *Business Process Management Journal* and the *Journal of Information Technology Education*. His latest research is centered on the monetization of enterprise knowledge management. Vik earned his undergraduate degree in business administration from Villanova University and a master's degree in information technology from Harvard University, where he received the Dean's List Academic Achievement Award.

Information Systems & Computing
Academic Professionals

STATEMENT OF PEER REVIEW INTEGRITY

All papers published in the Journal of Information Systems Education have undergone rigorous peer review. This includes an initial editor screening and double-blind refereeing by three or more expert referees.

ISSN 1055-3096