CMP
United Business Media

October
2007

BUSINESS
COMMUNICATIONS
REVIEW®

Securing The Campus: A Case Study

Jeff Dorsz

Here's how one community college district is protecting its network and data assets.

Jeff Dorsz, CISSP, manages information security and telecommunications for the South Orange County Community College District (www.socccd.org) and teaches at Saddleback College in Mission Viejo, CA. He has held managerial and senior level technical positions in security, network, systems and database administration, and has been a contributing author and technical reviewer for several security books. Jeff can be reached at jdorsz@socccd.org

It's not often that K-12 school districts or, in our case, a community college district, takes the lead in addressing an emerging network security issue. But educational institutions like mine are ahead of many commercial enterprises in facing the "dissolving perimeter" security problem, and in solving it with identity-based access and network policy enforcement.

As you may know, the "dissolving perimeter" problem refers to the inability of traditional edge devices, such as firewalls, VPN gateways and intrusion detection systems (IDS), to fully protect the assets on the internal network. This problem has arisen over the past decade as organizations allowed guests, contractors, business partners, customers, remote workers or, in our case, students, onto their enterprise network. Laptop computers add to the problem when users take them home or into Internet cafes, exposing them to all kinds of malware, then bring them back to their desks and plug them into the enterprise network.

Despite the security threats, however, users increasingly expect Web-based access to enterprise applications and data repositories whenever and wherever they can get "on the Internet." In short, the enterprise network has outgrown the old security paradigm, in which network access can simply be refused to "untrusted" outsiders and allowed to the "trusted" internal users.

Instead of an array of edge devices in the DMZ, a new approach is required that secures users and assets based on who they are and what they do, rather than on how they connect. Perhaps our experience in confronting this challenge and in rolling out our solution will offer some encouragement and insight to other organizations as they begin thinking about similar issues.

Facing The Challenge

My organization, the South Orange County Community College District (SOCCCD), located in Orange County, CA supports more than 38,000 students and 2,500 faculty and staff spread over two main campuses, including Saddleback College in Mission Viejo, and Irvine Valley College in Irvine, as well as the district office. Like other educational institutions, we experience the challenge of securing the network while giving appropriate access and services to a large student body that is using endpoint systems which we neither own nor control. The large student population, with students coming and going every semester, itself presents an administrative challenge in assigning and managing access rights to the users and their systems.

When I was hired in 2006, it was largely to help evaluate and address these security needs. The way had been paved, as the district administration already appreciated not only the importance of meeting these needs, but also that time and money would have to be spent.

We began with the overall goal of building a lasting security framework that could grow with the district and that would allow us to offer student and faculty access to the district information resources and services we had in mind. Designing a network security model to meet that goal in this kind of environment generally requires breaking down the security and risk management objectives into a series of specific projects, so we began by evaluating security initiatives. These included data-center access control, network access control, user compliance monitoring and internal intrusion protection (IPS) capability.

We also knew that, due to the large and distributed user community, there would have to be several phases of solutions rolled out. Our top priority was to secure the personal data (e.g., student data, and faculty and staff personnel records) stored at the district datacenter.

Our access policies are based on roles—student, faculty, administrator—and on other identity information

Adding Requirements

To secure the datacenter, we started off by evaluating the traditional approach of using LAN segmentation and firewalls to build an interior perimeter around the datacenter assets. This approach wouldn't work for us, however, mainly because our access control policies would be almost impossible to enforce using existing firewall rules.

Our policies are built around role-based definitions, such as "student," "faculty" and "administrator." Firewalls don't have any notion of user identity, nor of their roles in the organization, so it's practically impossible to map policy definitions into firewall rules.

Moreover, authorized users come from many different locations, and they frequently don't have static IP addresses or other machine identifiers that can be used to make an access policy decision. This is complicated by the fact that we don't own or manage the client endpoints.

We knew that the best solution would have to incorporate identity into the access policy and enforcement, and that, ideally, it would be integrated with our LDAP and Active Directory system which houses the student and faculty user data, including the roles that would determine access rights. After a few weeks of researching these requirements, the initial datacenter protection project evolved to also include network access control (NAC).

Specifically, our NAC requirement was to deny network access to users who did not have the proper credentials, and whose systems were not configured with anti-virus and operating system patch levels that meet our compliance and configuration requirements. This would allow us to put some defensive measures in place to contain the spread of malware, as well as restrict access to known users, and potentially assign them to a restricted use or quarantine VLAN accordingly.

Our existing network infrastructure already supported 802.1x for user authentication, so that requirement was added to the new NAC functionality. Finally, we also realized that, in addition to controlling user access and enforcing our identity-based network policies, we also would need to log and monitor user activity to detect any violations once users were allowed on the network.

Evaluating Our Choices

By January, 2007, SOCCCD had completed the requirements phase of the solution and had incorporated the datacenter access control, network access control and user monitoring projects into a single identity-based security initiative. Current NAC solutions provide a reasonable amount of identity tracking information and, once it's collected, it can be extended and integrated into an identity-based network policy control and enforcement solution. Unfortunately, we were under a tight budget for this initiative during the 2006-2007 fiscal year, and it looked like it would be difficult to implement both the NAC and the policy enforcement aspects of the project with a budget that was going to be less than $100,000.

For enforcement, we looked at highly scalable firewalls that would be able to accommodate the high datacenter bandwidth requirements, but they proved to be very expensive and none that we looked at could incorporate our identity requirements. Similarly, we found the NAC solutions we looked at were expensive based on the number of users and clients we were talking about, and many of them could not provide the user monitoring and access enforcement we required.

Then we turned to evaluating some other solutions. After looking at some of the stateful and application-level firewalls, we narrowed our scope down to a small number of NAC vendors that also claimed to provide in-line access control enforcement and both pre- and post-connect user checks and controls. In the first quarter of 2007, we brought in a selection of possible technology providers for evaluation and comparison. Of these vendors, we ended up choosing Nevis Networks

Figure 1 SOCCCD NAC Phased Deployment

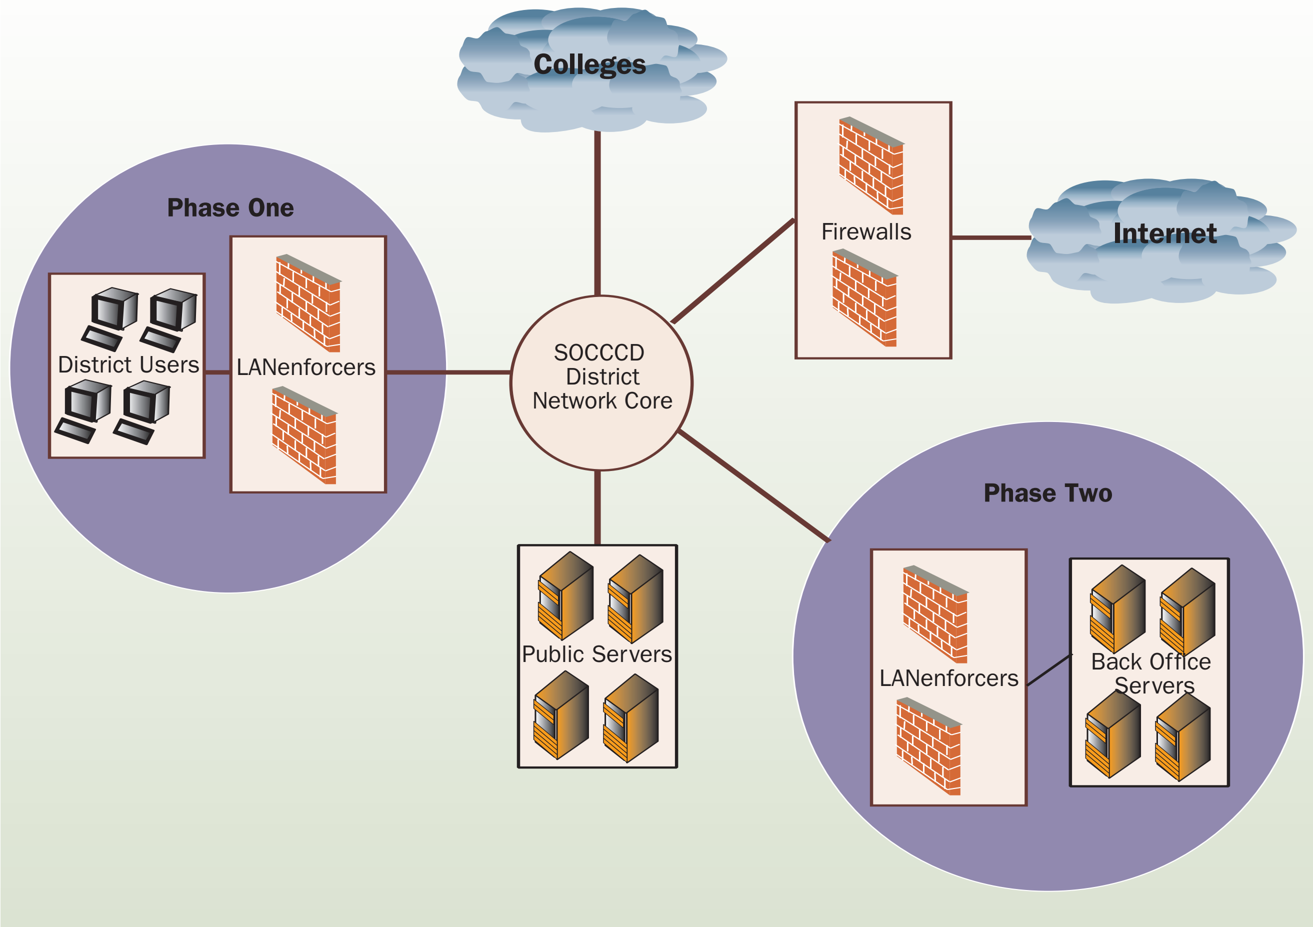

and their LANenforcer security appliance for deployment at the district office.

We liked the fact that the Nevis solution is built with identity-based policy enforcement in mind, so that it could take our student and faculty data from our Active Directory system, and so we could build role-based access control rules that would be enforced by the Nevis appliance, as we originally had hoped the internal firewalls would. Additionally, Nevis included both a pre-and post-connect NAC capability that transparently authenticates users to the network, scans the user's system for health and compliance purposes and assigns the user to a proper group for determining their access privileges.

The monitoring capabilities also met our requirements. Nevis analyzes packets to detect malware propagation, such as worms and Trojans, much like an intrusion protection system (IPS) would. We originally had in mind an IPS project for the following year, but Nevis could combine the solutions into a multi-pronged defense for about half of what we originally had thought we might have to spend on implementing the internal firewall approach.

Phased Rollout

By the end of March 2007, we had begun phase one of the deployment. We deployed the appliance in the district office, and applied the access control policies to the district office personnel only. This did include all our user roles, enabling us to fully build our security policies, so that later we could simply add more users with the various roles from the two campuses, while also familiarizing us with the new technology on a controlled scale.

After the initial success of the district office deployment and some testing in a live environment, we have now begun phase two of the deployment, which is intended to address our major concern, protecting access to the datacenter in the district office. We have deployed multiple LANenforcers and aggregated links in the datacenter, in order to secure specific back office server switches from the district network core. These two phases are shown in Figure 1, and we are also inserting the appliances into other district LAN segments, in order to extend the datacenter protection initiative to all district users.

Determining where to position these in-line appliances is the key issue when deploying a solution such as this across a distributed campus facility. Regardless of your configuration and rollout phase strategy, there are essentially two places that must be secured: Network choke-points, where traffic can't circumvent the system, and network edges, where the users connect.

If the primary objective is to authorize and secure network access, you'll probably secure the edges first, similar to phase one of our project. Otherwise, you'll probably deploy appliances at network choke points, as we did in phase two.

An important point for us was that this deployment didn't require a forklift upgrade of our network infrastructure. We can deploy in phases, in areas where we want to enforce specific network access policies or with specific user populations in mind. We also have been pleased to discover some cost savings and opportunities for business process improvement stemming from our project, which we hadn't anticipated.

We expect our solution will end up saving money that we otherwise would have spent on other types of devices such as firewalls and IDS/IPS. As to the process improvement, today, the user's role determines his/her district network access policy and the automatic updates let us enforce access controls without manually updating various system or network configurations.

Our solution will save us money and help to automate what used to be a time-consuming manual process

Conclusion

As we complete our rollout and tally up the costs and savings, and as the process improvements have a chance to mature and we learn to make maximum use of them, we think we are on the way to a very successful project. We knew we were taking on a common challenge that many organizations are facing with the dissolving perimeter problem, but we didn't know this would lead us to rethinking so much of what we knew about building a secure network infrastructure.

Perimeter-based solutions really are giving way to a new security model. In this emerging model, the entire network, in a sense, becomes a DMZ—protecting itself from all the potentially malicious users and unmanaged systems that, regardless of their threat potential, still have to have access to the internal enterprise network.

Fundamental to securing the network in this way is the notion of identity, which underlies all policy management and enforcement mechanisms. Identity-based features are available in many network access control solutions today. These capabilities associate users with their machine identifiers as they login, and monitor each user's session traffic thereafter, so any problems are clearly identified and tracked to the user causing them.

The business process improvements come from being able to translate well-defined identity-based policies right into the network infrastructure without the manual overhead of additional data entry or segmenting or reconfiguring the network, and from making it easy to manage users, not machine identifiers. It's a trend we think will become prevalent in the not too distant future□

Companies Mentioned In This Article

Nevis Networks (www.nevisnetworks.com)

Southern Orange County Community College District (www.socccd.org)

(#15459) Reprinted with permission from the October 2007 issue of Business Communications Review. Copyright 2007 CMP Media LLC.
For more information about reprints from Business Communications Review, please contact PARS International Corp. at 212-221-9595.