



Security Focus Shifting from Theft to Business Continuity

By Steve House

Disruptions Now the Most Costly

According to the Computer Security Institute's, "CSI/FBI 2004 Computer Crime and Security Survey", for the first time ever viruses became the most expensive security threat, followed by Denial of Service (DoS) attacks. These two categories passed Theft of Proprietary Information, which had been the most costly for the last five consecutive years. In fact, in the last year, viruses have caused nearly five times the expense of Theft of Proprietary Information, and the viruses and Denial of Service attacks together made up more than all eleven other categories combined.

Everyone realizes the obvious costs of virus infections. IT runs around fixing compromised servers and hosts, some systems might be brought down and some data may be lost. Another less obvious, but often more costly problem is the disruption to productivity when an entire site, or even an entire company, is slowed down or completely stopped due to the traffic generated when a virus or worm floods a network.

This flood of traffic has two potentially serious negative effects. The first is that the worm storm effectively functions as a DoS attack. The sheer number of new flows per minute that are generated from a single host can be enormous. Multiply that by potentially hundreds of infected hosts and this storm can mean disaster for anything in its path. This includes not only servers, but also any intermediate devices such as firewalls or routers. If a worm storm takes down a router, that's a serious problem because you've just lost that site.

The other nasty problem is the amount of bandwidth that these attacks often consume. If a remote site has a 128 Kbps link it could be swallowed up in no time by just a few infected machines. If the entire pipe is full of malicious traffic, the important applications are not able to function, and user productivity suffers. This impact to business continuity caused by worms, viruses, and DoS attacks has driven the cost so much higher than it previously has been.

Additionally, the frequency of these attacks has been accelerating, and each one is more sophisticated and complex than the last one. According to a July 2004 study from anti-virus firm Sophos, virus activity is up 21% in the first six months of 2004 compared to the same period in 2003.

Protecting the Perimeter is Only Part of the Solution

As the prevalence and maliciousness of worm, virus and DoS attacks has increased, the reaction has been predictable. Enterprises everywhere are focusing their efforts on preventing the next big attack. In fact, the overwhelming majority of IT security budget is spent keeping threats and

exploits out of the network. A recent Yankee Group study classifies this as Perimeter Security. Solutions in this bucket are the common security elements such as firewalls, anti-virus, URL filtering, Intrusion Detection Systems, and secure VPNs.

Does continued investment in perimeter security reach a point of diminishing returns? I would argue that it does. For example, if it takes \$10,000 to block 90% of intrusions, it might take \$100,000 to secure 99%, and \$1,000,000 to secure 99.9%. In other words the last 0.9% might take \$900,000, and nothing's ever 100% secure. The problem is that no matter how much money is spent deploying new technology, defining and enforcing strict corporate policies and training employees, outbreaks continue to occur.

Even if no laptops are allowed on the network, and all of the servers and desktops are regularly patched and updated with the latest hot fixes and anti-virus code, somehow the 'bad stuff' will get in. Whether it is an employee, partner, or contractor who breaks a policy and brings an infected laptop onto the network, or a new 'zero-day' exploit that takes advantage of a vulnerability that has just been discovered and doesn't yet have a patch available, it will happen. Attempting to keep them out is a noble effort and a battle worth fighting, but achieving one hundred percent success has become more of a pipe dream than a realistic goal. If the anti-virus or OS vendor doesn't even know about the vulnerability, how can you expect to protect yourself from it?

Maintaining Network Integrity Also Critical

The other component that is often neglected is what to do once the bad traffic enters the network. The Yankee Group, among others, is beginning to define this critical component as Network Integrity. Network integrity is about maintaining business continuity when a network breach does occur. This category includes, among other things, DoS protection, bandwidth management, and worm storm mitigation solutions.

Proactive enterprises have started to prepare for the inevitable. Understanding that every once in a while 'bad' things will happen, a plan and tools are put in place to minimize the impact they have on business productivity.

The first step to successfully weathering a network integrity breach is to be able to quickly—and automatically—identify that an anomaly is occurring on the network. Waiting for help desk calls reporting slow performance doesn't cut it. By then it is too late. Something needs to be continually monitoring the bandwidth consumed, new connections started, and application performance. Any abnormal change in these parameters should generate an immediate alert to the team so that they know to take a closer look before things get out of control.

The next step after realizing an exploit is taking place is to contain the impact. This means that the infected hosts that are causing the traffic are identified and 'dealt with'. Dealing with this traffic could mean taking infected hosts off the network, blocking traffic coming from them, or limiting the amount of traffic they are able to propagate across the WAN.

Limiting the traffic, specifically the number of new flows starting, allows all of the network devices such as routers and firewalls to avoid overload conditions. Contrary to the common perception that devices are limited by the bandwidth they can support, new flows per minute is much more of a limiting factor. For example, a firewall might easily manage one FTP flow consuming 50 Mbps, but could crumble under thousands of tiny new flows that are typical of viruses that may only generate a few Mbps of traffic.

An additional benefit of limiting this traffic, is that more bandwidth is available to all of the business applications that need to run over the already bandwidth constrained WAN links. In other words, reducing virus traffic crossing the WAN allows all of the applications to perform better.

Finally, and most importantly, critical applications should be protected at all times. Using the common highway analogy, if you have an express lane that only important vehicles are allowed access to, it doesn't matter if the rest of the highway is crammed so full of cars that traffic grinds to a halt. The vehicles in the reserved lane continue to flow quickly. The same goes for application traffic over WAN connections. Guaranteeing the critical applications get the bandwidth they require for peak performance applies just as much if the rest of the traffic is from virus or DoS attacks as it does if it is from large e-mail attachments, casual web browsing, and file sharing.

Traffic Management Solutions Complement Perimeter Security

Traffic management solutions provide the tools to do just what the name implies—they identify anomalies, contain bad traffic and protect good traffic.

The good news is thousands of enterprises have already deployed these devices on their network for network monitoring to understand what applications are running on the network, bandwidth management to provide quality of service to critical applications and compression to get more data through existing WAN links.

The ROI for these appliances has been easily justified by the other functions they are performing on the network, and therefore using them to help maintain network integrity is almost a freebie.

The problem is that many of these organizations haven't yet realized the power these appliances have to help in their fight to maintain network integrity. In many cases they have been deployed for just one or two purposes without ever going back and thinking about the additional value they can provide. They do that one job, and they do it well—why go back and mess with them?


Existing customers of the application traffic management devices really need to explore the possibilities and realize all of the benefits they can bring to the table. It's like buying a car and only driving it to work, and not thinking about all of the other cool places it could take you.

Organizations that haven't yet deployed application traffic management solutions should be looking into them. When a single tool can provide granular visibility into how a network and the applications running over it are performing, ensure critical applications get the performance they need, double or triple the effective bandwidth with compression, AND help survive worm, virus and DoS attacks, the resulting ROI can often be

measured in months or even days. In the case of surviving an attack, it will normally pay for itself the first time the network would have gone down, but didn't.

Perimeter Security and Beyond

It is important to think about perimeter security as the first line of defense, but not the only component in your fight against worms, viruses and DoS attacks. Assuming perimeter security will be 100% successful is a recipe for disaster. Instead, network integrity solutions should also be deployed to mitigate the business impact caused by a flood of malicious traffic that accompanies these attacks.

Application traffic management solutions provide the tools necessary to maintain network integrity, nicely complementing the efforts spent on perimeter security. Having both these solutions in play provides the extra level of confidence that even if something bad does happen; the business will continue to run effectively. Not having them could put you in some awkward situations that may end up costing even more—your job. 

Steven House is a senior manager of product management at Packateer, Inc., where he has helped drive the security features of the company's Application Traffic Management system to enable it to classify and control malicious traffic traveling over WAN. Steve has also held product management roles at CasheFlow (now Blue Coat Systems) and Newbridge Networks. Steve earned his degree in Bachelor of Science in Electrical Engineering at Duke University.