



# Global Information Assurance Certification Paper

Copyright SANS Institute  
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

## Interested in learning more?

Check out the list of upcoming events offering  
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"  
at <http://www.giac.org/registration/gsec>

**Abstract:** A Threat Assessment and Cost-Benefit Analysis of Deployment of Host Firewall Software in Medium to Large Networks

Research on the security posture of corporations suggests that host perimeter protection is routinely used for protection against viruses, but rarely used against the threat of intrusion or other malicious network traffic. This is probably due to a misconception on the part of systems and network administrators regard the threat of penetration of their firewall device, the presence of real threats of intrusion in the wild, or the cost of installing and maintaining such software. A careful analysis (including looking at survey on actual security incidents) reveals that networks with firewalls are still vulnerable, and that exploits exist and are being used to penetrate even firewalled networks.

Given the real threat and the high stakes, only software that works well in a networked and centrally administered setting can be deployed in a cost effective manner. Looking at the features such “host firewall software” would have can, in the longer term, provide a roadmap for a development of next-generation security tools. In the meantime, it provides a template against which features of existing personal firewall software can be measured for deployment in such an environment. Though each corporation and networked organization needs to perform its own threat assessment and cost-benefit analysis to decide whether to employ such software, it is clear that some software exists today with many (if not all) of the required features, and the security posture of many corporate networks could be significantly enhanced by the carefully planned deployment of such tools.

## **A Threat Assessment and Cost-Benefit Analysis of Deployment of Host Firewall Software in Medium to Large Networks**

To provide the layering and redundancy necessary for optimal security, a network security administrator needs to configure an appropriate combination of security measures both at the network border and at the perimeter of internal hosts. A firewall is the obvious tool for the network border; there are more choices for the host perimeter. Host-based intrusion detection software (or HIDS) is usually designed for deployment on a small number of servers and other critical hosts in any network, and it is often impractical and expensive to deploy such software more generally. Remaining choices to protect the perimeter of the "non-critical" host (which is to say, most hosts on the network) include anti-virus software and personal firewall software. Though serving somewhat different functions, these two tools provide protection at the same location in the network topology; patterns of deployment of these two in real-life networks, however, diverge significantly. Anti-virus software has become nearly universally deployed, at least in major corporate networks, while personal firewall software is deployed much more rarely in such an environment. The likely reasons for this difference include perceptions (often mistaken) of risks and the protection provided by each, and perceptions (often correct) of the more difficult task of effective deployment of personal firewall software in a typical corporate network. In this paper, I will outline statistics and strategies for raising awareness of the real threats which can be addressed by effectively deployment of personal firewall software in networked environments. Still, recognizing the need is only half the battle, so I will also describe the features needed for a next generation of "client firewall software" optimized for such environments, and examine a couple of current personal firewall products to see how well they match up against the features such a next-generation product might have.

Models of defense-in-depth network security are based on the concept that no single layer of protection can provide sufficient protection between any threat and any target. Put another way, defense-in-depth is proactive security. The systems administrator using this strategy, for example, doesn't need to read about exploits that can penetrate the network perimeter (or find himself on the receiving end of such an exploit), to decide to react with an additional internal layer of defense, but rather always looks for those locations in his network where the fewest layers of protection exist, and searches for ways to strengthen his defenses there. There is an interesting comparison, then, to be made between the deployment of various tools and strategies for protecting against viruses on the one hand, and on the other, protecting against trojans and penetration attempts. Though each of these threats target individual network hosts, in each case, some measures can be taken at the network perimeter to protect against such threats, though such measures provide imperfect protection. And in each case, other measures, also imperfect, exist to provide another layer of protection at the perimeter of the individual host. Given all of this information, one would imagine that a defense-in-depth approach would suggest a multi-layered approach to each of these threats.

Data exist, however, to reveal that actual patterns of deployment of security tools tend not to follow this approach, suggesting either that defense-in-depth is not

(yet) the standard strategic approach to security, or else that some external factor is an obstacle to such an approach, such as misperceptions of the threats, or high costs of deploying certain strategies. Specifically, according to the *2002 CSI/FBI Computer Crime and Security Survey*, 89% of respondents report using a firewall as a security tool, and 90% report using anti-virus software, but only 60% report using intrusion detection software.<sup>1</sup> (There is no category at all in this report for personal firewall software or any other category of host-based security tools except intrusion detection. Possibly firewall software is used in these contexts so rarely as to fail to warrant inclusion. In any event, it seems likely that respondents would consider it a form of intrusion detection software, if they took it into consideration at all in their responses.) We can conclude at least that antivirus software has reach close to universal implementation,<sup>2</sup> while personal firewall software has not. We can also infer that anti-virus software is the only host-based protection on most networks; that is, if we assume that, in most cases, respondents who says they use intrusion detection software are using convention HIDS on a few keys hosts, then it seems likely that most hosts on most networks have no protection except antivirus software at their perimeter. Or, in other words, many networks provide no defense between the firewall and most of their hosts against threats which cannot be handled by their anti-virus software<sup>3</sup>.

This asymmetry in security software deployment is curious. Three possible explanations for these differences seem possible, and need to be considered.

- Maybe system administrators making decisions about deployment of security tools feel that a firewall protects sufficiently at the network perimeter against intrusion and malicious software to eliminate (or at least greatly reduce) the need for additional protection at the host perimeter. Presumably, and by contrast, they feel that any network perimeter protection against viruses (such as protection at the mail server) is not sufficient to eliminate the need for host-based protection against this threat.
- Or, maybe they do not believe that the threats against which personal firewall software protects are merely theoretical, as opposed to the widely publicized threat of viruses, with its well-known and high costs.
- Or, finally, maybe these system administrators don't feel that the protection afforded by this software outweighs the challenges and difficulties of installing and maintaining it; by contrast, antivirus software is relatively easy to install, configure and maintain throughout a network.

Probably all three of these factors play a part, in different degrees. Each point can be countered, however, and taken together, these counterarguments add up to a strong case for wide deployment of some form of personal firewall software in corporate and other networks, and point toward some best practices that make such a deployment feasible.

*Point 1: Use of a firewall does not eliminate the need for additional host-based protection*

Simply using the same term to describe a hardware firewall and personal firewall software implies that there is some redundancy between using both, and therefore that the presence of a hardware firewall eliminates the need for such software. Even if the principles of defense-in-depth didn't dictate the creation of such multi-layered protection, analysis of known exploits and vulnerabilities of firewall would. For example, search of the term "firewall" on the cert.org list of vulnerability notes returns 142 results.<sup>4</sup> While not every returned result represents a vulnerability that could permit a firewall to be bypassed, many do, the following are a few of the more serious vulnerabilities taken only from the first page of the search results returned:

- "A somewhat common configuration of Cisco PIX firewalls may permit a window of opportunity in which an intruder can bypass the firewall."<sup>5</sup>
- "The 'netfilter' firewall subsystem included with Linux kernel versions 2.4.x contains a vulnerability that may allow remote attackers to reach hosts that should be protected."<sup>6</sup>
- "If any rules include the 'Fast Mode' option, Check Point Firewall-1 and VPN-1 will incorrectly allow unauthorized connection attempts to hosts that should be restricted."<sup>7</sup>

And of course, these are examples demonstrating only that an apparently properly configured firewall can, in some cases, be compromised so as to allow in or out traffic that should be stopped. Even if it were not for vulnerabilities such as these, there are other scenarios where a firewall alone provides insufficient protection. A firewall can be circumvented via dialup to an internal host, or through a VPN connection with an insecure password. Even a firewall whose integrity remains uncompromised is only as strong as the policies it runs, and corporate pressures may exist toward more permissive policies which could in fact allow malicious traffic to pass alongside legitimate traffic.

Firewall software does not eliminate the need to address all of these concerns. However, neither can even the most effective effort to address all of these concerns eliminate the need for firewall software, to act as a second line of defense. Fred Langa of InformationWeek.com summarizes effectively: "[I]t's risky--almost foolish--to depend on a single line of defense. . . . [A] firewall can fail; no piece of hardware or software is perfect."<sup>8</sup> If the firewall fails in any way, software protecting the host perimeter against a hack becomes the critical last line of defense, the one that may ultimately protect critical information or resources, or may simply guarantee that a hacker cannot quietly compromise an internal host. Consider here again the analogy to protection against viruses. In a network where virus protection exists only at the network perimeter, we can imagine a virus somehow bypassing this protection and infecting a host. The network perimeter protection would, however, continue to protect against other infections. The same is *not* true with intrusion protection; once a single host is compromised, a hacker can often use this machine to send traffic unnoticed through the firewall, and of course traffic between this compromised internal host and its peers will not travel through the firewall at all. The intrusion protection at the network perimeter is essentially removed from the equation now,

and protection at the host perimeter may be the only protection against the effective escalation of this attack.

This leads directly to the second point that Fred Langa makes, as he goes on to point out that

a conventional firewall may do nothing at all to protect against attacks that originate on the "safe" side of the connection or that attempt their dirty work via the usually lightly guarded outbound Internet link. These attacks can result from intramural hacking (across the local network) or from Trojans, worms, and "phone-home" spyware installed on local systems.<sup>9</sup>

In other words, not only may a firewall be compromised or circumvented, but it may be entirely irrelevant in protecting against attacks launched from the internal network (either from an compromised internal host, or perhaps from an authorized user on the inside), or against unauthorized outbound traffic. In these cases, host-based protection wouldn't be an additional layer of defense, but would potentially be the *only* layer of defense. It is hard to imagine, from the perspective of defense-in-depth, or really from any perspective that takes security seriously, a strong case to be made against deploying a tool to protect against a threat that would otherwise be largely or entirely unaddressed by the resulting security posture.

*Point 2: The threat from intrusion and trojans is real, and the costs associated with it are significant.*

The empirical evidence that these threats are real is well documented. Recall from above that in the *2002 CSI/FBI Computer Crime and Security Survey*, 89% of respondents reported using a firewall as part of their security configuration. Add to that the finding that 40% of respondents reported that they experienced one or more security incidents involving penetration of their network from outside, twice the rate of just five years ago.<sup>10</sup> (In fact, penetration is one of only two of the twelve categories of attack or misuse that show a clear upward trend over the last five to six years, the other being denial of service.) These attacks against corporate networks are occurring, and their numbers are growing. Though they may not reach most such networks, they reach many. And, given the 60% rate of intrusion detection software, it is possible that there is a great deal of unnoticed intrusion and network penetration beyond that reported here. Perhaps most critically, however, notice that the number of respondents reporting such an attack is far greater than the number of respondents not using a firewall. In other words, most of these attacks from the outside which achieved network penetration occurred *despite* the presence of a firewall. Not only are these attacks real, but they are occurring even when a firewall is in place.

If the threat is real, however, the publicity behind it is not, at least when compared to the virus threat. Corporations are notoriously reticent to publicize their security incidents, perhaps with good reason. But world-wide virus incidents are visible to everyone in a way that individual network penetrations (however many) are not. Consider the LoveLetter virus, which appeared in May of 2000, and which one

article documented as costing companies nearly a billion dollars in clean-up, and another \$7.7 billion in lost productivity.<sup>11</sup> The critical point here is not the specific dollar figure, but the fact that the nature of virus outbreaks allows the cost to be so well and widely documented. Large scale virus outbreaks are highly visible, affecting many corporations and networks at once, and the mainstream media have established a habit of reporting such outbreaks. This wide scale of publicity, and hence awareness, is probably part of the reason that the *2002 CSI /FBI Computer Crime and Security Survey* reports a drop in respondents experience virus attacks, from 94% to 85%,<sup>12</sup> despite the fact that the current year's data included the time period of the outbreaks of the SirCam, Nimda and Code Red viruses. In one news story written during the SirCam outbreak, "anti-virus experts" explicitly attribute the lower number of infections to more effective deployment of anti-virus software by corporations that had realized the danger and costs of earlier outbreaks.<sup>13</sup>

By contrast, the threat of network intrusion is not so well publicized. Incidents tend to take the form of individual attacks against a corporation and its network. Even if such an attack succeeds, it may be entirely invisible to the outside world (web defacement attacks being one of a few noticeable exceptions), and may go unreported by the victim. In other words, the relative lack of publicity and awareness regarding such intrusion attacks may very well be a result of the individual nature of those attacks, rather than their rarity. The consequence is clear: Without hard evidence to indicate that this threat is both real and common, even expert systems administrators can take away the wrong impression. And that wrong impression can have direct consequences in the security posture.

Perhaps the most important evidence to persuade them (and those above them in the chain of command) is the cost. Turning once again to the *2002 CSI /FBI Computer Crime and Security Survey*, we see that the average annual loss to respondents reporting quantifiable losses due to network penetration is \$226,000, slightly below the \$283,000 in analogous losses due to viruses. In each of the five previous years the survey has been conducted, they costs associated with intrusion incidents has been the higher of the two. The conclusion to be reached is simple enough: the threat is real, and the stakes are high. If the lack of host-based intrusion protection software relatively to comparable anti-virus software results from a perception that the threat can be safely ignored, the evidence demonstrates that this perception is false.

*A careful analysis of the costs and benefits of deploying personal firewall software on corporate networks indicates it is usually worthwhile to do so.*

On this point, more than the previous two, the justification for failing to deploy personal firewall software may have some real validity. Given the state of personal firewall software on the market, and the challenges that arise in attempting to deploying and effectively manage this software in a medium to large-sized network, it is not immediately obvious that the benefits outweigh the very real costs in limited time and effort. Much of the current generation of personal firewall software is very effective as a *personal* security tool; that is to say, it is very effective on personal stand-alone computers or on small and simple home networks. This is exactly what

we should expect, since this is the intended audience for the software. But this also will mean that the same software will tend to be poorly designed for a networked setting.

We need to consider exactly what sort of features would make personal firewall software effective and affordable in a typical medium to large TCP/IP network with Windows-based clients and a small number of administrators responsible for security and other administrative tasks. We can safely assume that the users of this software will be IT professionals with a good working understanding of TCP/IP, and of firewall policy as it is already implemented at the perimeter of their network. As the list of features required below makes apparent, the software that fits the bill isn't really *personal* firewall software at all; I'll consider it a separate class of software, which I'll refer to as "host firewall software." Such software would, optimally, have the following features to guarantee adequacy of the protection it provides:

1. Powerful rule configuration: To work in conjunction with the security policies in place on the network firewall, the software would have to support filtering rules based (at least) on application, remote IP address, local and remote port numbers and protocol.
2. Configurable alerts: Key events need to be able to be defined which trigger email or other sorts of urgent notifications of administrators.
3. Persistence of protection: This software has to be difficult for a local user to disable or shut down, and there should be a mechanism for alerting an administrator if it is not running on an active computer where it should be.
4. Shielding from the user: Local system users should not be able to modify policy locally, or override it to allow access applications greater permission than the rules permit.

Additionally, the following features would be required to guarantee the affordability (in terms of both financial and human resources) or the acquisition, deployment and maintenance of such software.

1. Central administration: This software would have tools to allow administrators to review, set and update policy on many workstations from a central location. It is essential that updates could occur either on-demand, or in some sort of automated roll-out or push to clients.
2. LAN readiness: This software will need to be easily configured to allow for the sorts of traffic common on LANs, but without doing so in a way that is either so permissive or inflexible as to undermine the goals of adequate protection against any possibly hostile internal traffic.
3. Central logging: This software must allow for some type of central logging of events, for ease of review on a regular basis.
4. Reasonable cost: This software should be comparable in price to that of anti-virus software, on the order of \$50 or less per client installation, and probably substantially less for bulk or site licensing. This is a purely pragmatic standard for pricing of course, on the logic that budgets have already been approved for security tools for these amounts within most organizations in the recent past.

Software with all of these features would be no more difficult (and probably substantially easier, in many cases) to deploy than typical anti-virus software (and, in fact, some of these features, including some of the centralized deployment and monitoring features, would do well to be adapted for enterprise deployment of anti-virus client software as well, though that goes well beyond the scope of this discussion). That is not to say that there wouldn't be any cost associated with such deployment. It would add a substantial task to the already full workload of administrators handling network security issues, but it would provide them with a valuable new tool as well, one that, given the costs and risks outlined above, would be well worth deploying – if software with all of these features was, in fact, available.

I have no doubt that, over time, the demand for software with these features will increase, and market forces will encourage development in these directions. There already some indication that software approximating some of these features is being developed and hitting the marketplace at this time, such as an enterprise products from ZoneLabs.<sup>14</sup> This product is designed and priced only for very large enterprise deployment, however, and really is not appropriate for deployment in networks any smaller.<sup>15</sup> The appearance in the market might be evidence of movement in the right direction, but this is neither the ideal tool for medium to large networks, nor is it even one that can be reasonably acquired and then adapted for such use. To find a tool that can serve that function, rather than examine an enterprise tool, and then look for ways to scale it down, we need to look for small personal firewall products, and look for ways to scale it up.

No discussion of specific personal firewall software products can avoid discussion of ZoneAlarm from ZoneLabs. Whether or not it is representative of all such software in the marketplace right now, it is at this point so widely deployed and well-known that its features comprise the full personal firewall experience for many users of such software. It is hard to get a quantitative measure of personal firewall software on home computers, but download statistics at CNet's download.com website reinforce this impression: in one week, the free version of ZoneAlarm had been downloaded over 180,000 times, and the "pro" version had been downloaded over 12,000 times. The closes competitor, Tiny Personal Firewall, had been downloaded just under 10,000 times.<sup>16</sup> Even considering that many packages may be disproportionately acquired by download directly from the manufacturer's website or as part of a bundle of other products, these overwhelming numbers demonstrate the predominance of this product.

This is important to the discussion at hand because ZoneAlarm's popularity is due the very features that make it so poorly suited for use in a networked setting (and in turn, it may create a popular impression among IT professionals that all such software is poorly suited for such a deployment). A careful consideration of the software against some of the criteria list above makes this pretty clear.<sup>17</sup> Regarding adequacy of protection, the "Pro" version of ZoneAlarm includes adequate rule configuration and shielding of the user at a reasonable price (defined by the "not costing much more than your anti-virus software" rule). However, ZoneAlarm does not allow for any sort of configurable alerts, or persistence of protection; the result is that critical security events could go unnoticed, or the software could be easily

disabled or uninstalled without an administrator's knowledge. Perhaps even more critically, however, are the ways in which ZoneAlarm fails to meet the host firewall software criteria for reasonable cost of use. The actual cost to purchase the software is reasonable, by this standard": the "Pro" version costs about \$50 per copy, but has substantial discounts for bulk licensing; annual licenses for support and upgrades costs about another \$20 per machine, regardless of volume. None of the other criteria for cost, however, are met. Though the software allows for definition of a LAN subnet with a separate (presumably more permissive) set of filtering rules, there is no automation of rules to allow Windows LAN traffic. Additionally, and most importantly, there is no mechanism for centralizing administration or policy roll-outs, nor is there any way to centralize logging. The administrator using this software would have to spend a great deal of time visiting individual workstations each time policy need to be set or changed, and review of logs presumably would happen rarely, if at all.

Remember, however, that my claim above was that deployment of personal firewall software in medium to large networks is usually justified by the protection it provides, as measured against the cost. Clearly ZoneAlarm can provide this protection at a sufficient cost (despite it being an excellent product for home users); but I will argue that other such software packages can, and will take, Kerio Personal Firewall software as an example. I am not claiming that this software is the only software, or even the best, for this purpose, but simply that it is an example that demonstrates that software exists which can be effectively deployed, while we all wait for true host firewall software to appear.

Consider the features of Kerio Personal Firewall against the protection standards given above for ideal host firewall software. As with ZoneAlarm Pro, Kerio provides powerful rule configuration and user shielding. On the other two criteria, while imperfect, Kerio has some important advantages over ZoneAlarm. Kerio does *not* have a mechanism for defining alerts, but since it does allow for logging to a syslog server, alerts could be configured at this server and delivered to administrators. Additionally, there is no automated featured providing persistence of protection, but since any installation of Kerio Personal Firewall can connect to any another (after providing configurable passwords), it is possible for an administrator to manually verify that the software is running on any particular workstation. If it is not, but the machine can otherwise be determined to be running (via a ping or any other method), then an unprotected workstation can be identified. Obviously, only allowing for such checking manually is still a substantial shortcoming, but it does provide some additional protection.

Similarly, Kerio rates reasonably well against the standards for cost of deployment and use given above. Kerio sells the software for about \$40 per installation, with substantial discounts for volume purchases. The software includes, as mentioned above, the capacity to log to a central syslog server, allowing for relatively easy review of all logs from one location. The software has options for permitting LAN traffic, which, when chosen, permits all port 137, 138 and 139 TCP traffic within the defined local subnet; however, additional rules *cannot* be defined to further restrict such traffic. Finally, Kerio does allow for some central administration: any installation can connect to and administer any other; appropriate configuration with password protection can essentially create a system where an administrator at

any desktop can review settings at any other installation. Additionally, settings on any machine can be saved to a file, and then locally or remotely deployed to any other. What *cannot* be done, however, is a single centralized and automated push of new settings to all installations. What this means is that a new configuration could be deployed on all workstations in a network centrally, but it would require manual updates on a host-by-host basis.<sup>18</sup> For very large networks, the time commitment of such updates would make it impossible. In smaller networks, however, careful deployment to minimize the frequency of updates, combined with effective use of staff resources (since it requires little IT expertise to actually update configurations, once they new rules are determined) can make such manual updates a worthwhile part of an overall security strategy.

Each organization needs to decide based on its own priorities and budgets if these considerations make it worthwhile to deploy Kerio or any other personal firewall software as a means of host perimeter protection, and probably some study and trials would be appropriate for any organization attempting to make this decision. However, the evidence at this time suggests that corporations are not routinely investigating such a strategy, and are not taking seriously the threat against which such software protects. My argument here is not that some narrowly defined course of action is necessary, but rather that there is a common misconception that an effective network security posture can be established without careful consideration of protection of the host perimeter against intrusion or other malicious traffic. The danger is real, the stakes are high, and the tools and techniques to provide protection are available; a serious defense-in-depth approach requires that such tools be used, or that their non-use be justified only by the most careful analysis of the cost and threat.

## References

Becker, David. "SirCam worm unlikely to become epidemic." C-Net News.com. 24 July 2001. URL: <http://news.com.com/2100-1001-270481.html> (27 July 2002).

Langa, Fred. "Langa Letter: Firewall Feedback." InformationWeek.com. 15 April 2002. URL: <http://www.informationweek.com/story/IWK20020412S0009> (27 July 2002).

Lemos, Robert. "Lessons of 'Love' virus still sinking in." C-Net News.com. 4 May 2001. URL: <http://news.com.com/2100-1001-257095.html> (27 July 2002).

Private phone conversations with various ZoneLabs sales and support staff in phone conversations, 22 July 2002 and 23 July 2002.

Power, Richard. "2002 CSI /FBI Computer Crime and Security Survey." Computer Security Issues and Trends Vol. VIII, No. 1, Spring 2002. URL: <http://www.gocsi.com/pdfs/fbi/FBI2002.pdf> (27 July 2002).

"Search", URL: <http://search.cert.org>. (27 July 2002).

"Security and encryption." URL: <http://download.com.com/3150-2092-0-1-5.html?> (12 August 2002).

"Vulnerability Note VU#230307: Linux kernel netfilter IRC DCC helper module creates overly permissive firewall rules." URL: <http://www.kb.cert.org/vuls/id/230307> (27 July 2002).

"Vulnerability Note VU#446689: Check Point FireWall-1 allows fragmented packets through firewall if Fast Mode is enabled." URL: <http://www.kb.cert.org/vuls/id/446689>. (27 July 2002).

"Vulnerability Note VU#6733: PIX 'established' and 'conduit' command may have unexpected interactions." URL: <http://www.kb.cert.org/vuls/id/6733> (27 July 2002).

"Zone Labs: Enterprise Solutions." URL: <http://www.zonelabs.com/store/content/company/corpsales/intOverview.jsp>. (15 August 2002)

"ZoneLabs Integrity Datasheet." URL: [http://download.zonelabs.com/bin/media/pdf/Int\\_data910.pdf](http://download.zonelabs.com/bin/media/pdf/Int_data910.pdf). (15 August 2002)

---

## Notes

<sup>1</sup> Power.

<sup>2</sup> It may seem that strong to claim that implementation is "close to universal" based on 90% reported deployment. In each of the four previous years this question was asked, however, the reported rate of deployment of antivirus software ranged from 96% to 100%, suggesting that the 90% rate is either a statistical anomaly, or else represents an unlike move *away* from antivirus software. Considering that the report shows a drop in the rate of deployment of 10 out of 11 categories of security technologies, it seems most likely that all the drops are the result of some broader change, such as a change in the profile of respondents.

<sup>3</sup> I recognize that here and throughout, I am arguing as though the distinction between anti-virus software and personal firewall software were a sharp one, when in fact it is not. Increasingly, in fact, viruses, worms, trojans, and other intrusion tools overlap greatly, so it is inevitable that the tools to defend against them will as well. However, the fact remains that there are distinct security tools available, each of which primarily serves to mitigate certain types of threats, and makes little or not effort to mitigate against others. For shorthand, I'll refer throughout to the threat against which personal firewall software (and firewall hardware devices) protect as "intrusion", understanding this term to encompass an array of related threats.

<sup>4</sup> "Search".

<sup>5</sup> "Vulnerability Note VU#6733.

<sup>6</sup> "Vulnerability Note VU#230307.

<sup>7</sup> "Vulnerability Note VU#446689.

<sup>8</sup> Langa.

<sup>9</sup> Langa.

<sup>10</sup> Power.

<sup>11</sup> Lemos.

<sup>12</sup> Power.

<sup>13</sup> Becker.

<sup>14</sup> See, for example, "Zone Labs: Enterprise Solutions" and "ZoneLabs Integrity Datasheet.", for information about ZoneLabs Integrity Product. A full review of this product is well beyond the scope of this paper, but since this product uses ZoneAlarm as the client at each desktop, much of the discussion of ZoneAlarm traffic filtering below applies here as well. Additionally, the (presumably) high cost of this product puts it well beyond the scope of this discussion on the effective use of inexpensive tools for network host perimeter security.

<sup>15</sup> Private phone conversations.

<sup>16</sup> "Security and encryption."

<sup>17</sup> This discussion is based specifically on an analysis of features of version 3.1 of ZoneAlarm and version 3.0 of ZoneAlarm Pro, but after working with several different recent versions, I have not seen substantial variation in the features discussed here.

<sup>18</sup> Perhaps a logon script could be written to copy the configuration to each machine, but my own experiments with such a script were not successful, because Kerio's software runs as a service loading the rules into memory, and ignoring rule changes does through file copying after the service is running. Additionally, not all machines are routinely logged onto, potentially delaying deployment of critical updates.