

**Beyond Detection:**

**Neutralizing Attacks Before They Reach the Firewall**

---

A Network Security White Paper from ForeScout Technologies

Summer 2002

## Executive Summary

The sophistication and volume of security threats facing businesses have grown significantly. To make matters worse, companies have become extremely dependent on their computing/communications infrastructure – making the potential impact of those threats greater than ever.

However, as daunting and as critical as the challenge of info-defense may be, security managers don't have infinite budgets. So they must somehow protect their organizations without placing undue strain on limited capital and human resources.

To protect the enterprise, security managers have deployed a variety of technologies. While these technologies are useful for defending corporate assets, they have limitations. Firewalls, for example, may be configured to block certain types of traffic, but attackers still find ways to exploit legitimate traffic types to mount their attacks.

Intrusion detection (ID) presents its own difficulties. ID systems only detect attacks that fit an established pattern or “signature.” This leaves the network vulnerable to new, undocumented attack strategies. ID systems also tend to yield a large number of false positives – thereby wasting staff time and eventually causing a real attack to be ignored. Other types of anomaly recognition systems are similarly prone to generating false positives, since they trigger alerts whether a deviation has an innocuous or a malicious cause. Ultimately, ID and anomaly systems are reactive rather than pro-active.

Some organizations have also deployed so-called “honeypots” to lure potential attackers away from the enterprise and document attack attempts. Honeypots, however, can't fully guarantee that they – rather than the enterprise network – will be the target of the next attack.

**ForeScout Technologies' innovative ActiveScout™ solution fully addresses the shortcomings of these conventional security tactics by effectively preventing attacks at the network perimeter – without requiring that the nature of those attacks be known in advance and without requiring security staff to constantly deal with attack alerts that may or may not be valid.**

ActiveScout delivers this unique protection by monitoring the reconnaissance activity that invariably precedes an attack. Rather than incessantly alerting security professionals about this activity, ActiveScout automatically responds to pre-attack recon with “deceptor” data, which effectively prevents any subsequent attack from reaching the corporate network. The network is thus protected against both known and unknown attacks.

With ActiveScout, information security managers can:

- **pre-emptively neutralize attackers before they even reach the firewall**
- **gain a new level of protection from both known and unknown attacks**
- **reduce the workload on their already overburdened staffs**

ActiveScout is therefore an indispensable security solution for organizations seeking maximum protection for their critical information resources. Its powerful, pro-active approach to the prevention of security breaches sets a new standard for perimeter defense – delivering unmatched effectiveness and resource-efficiency.

## The State of Perimeter Security

The network perimeter has become a primary area of concern for corporate information security managers. Because networks are now invariably connected to one or more outside networks – including, of course, the Internet – a wide range of threats now face the corporation. These threats are exacerbated by the fact that internal systems are actually quite vulnerable to all kinds of exploits. Plus, the widespread availability of recon tools has made it easier than ever for even novice attackers to nibble away at enterprise security from the network perimeter inward. So security professionals are under a lot of pressure to prevent any penetration of that perimeter.

Unfortunately, corporate infosec teams have very limited resources. There is not an infinite supply of experienced, skilled security professionals. And even the most skilled professional can only perform a limited number of security tasks at any given time. Combine these human limitations with limited security budgets and it's clear that corporate infosec teams must be as concerned with being efficient as they are with being effective.

Security professionals have focused on multiple technologies to keep their networks safe: firewalls, intrusion detection (ID), and honeypots.

### *Firewalls*

Firewalls are usually the first component of any perimeter defense. Firewalls perform the critical task of filtering traffic crossing the network boundary. This filtering is done according to predefined security policies, which can be specified at the network layer and/or at the application layer. For instance, the RealAudio™ Streaming Protocol might be handled at the network layer, while SMTP may be analyzed at the application layer to search e-mail attachments for virii and other “malware.”

Firewalls, however, do not provide airtight perimeter protection. After all, they have to let legitimate traffic through. **Their main deficiency stems from the fact that they utilize static, manually configured policies to differentiate legitimate traffic from non-legitimate traffic.** Those policies can vary in effectiveness, depending on the expertise of the security manager and the complexity of the network environment.

Also, once a static policy is defined, the firewall can't react to a network attack – nor can it initiate effective counter-measures. So if policy makes a certain network service available, it will remain available even if that service is used to mount an attack. **In other words, firewalls may be strong – but they can't respond to security incidents as they occur.**

### *Intrusion Detection*

A second layer in the perimeter defense is network-based intrusion detection (ID). Theoretically, network-based ID systems work like a burglar alarm, alerting infosec managers if an attack is taking place so that they can respond accordingly. They trigger these alerts by detecting anomalous traffic patterns or “signatures” that are characteristic of an attack.

However, today's ID systems exhibit several shortcomings that limit their usefulness in protecting the network.

**The first of these shortcomings is their propensity to generate “false positives.”** That is, they frequently generate alerts about an attack when none is taking place. False positives, like all false alarms, create two problems: 1) they waste the valuable time of staff security managers, and 2) they create a “Cry Wolf” environment where real attacks may be ignored.

In fact, when initially installed, it is not uncommon for an ID system to have more than 95% of its alerts turn out to be false positives. This hypersensitivity can be reduced by “tuning down” the system and making it more selective. But ultimately, the only way to eliminate false positives would be to tune the system down to the point where it would also ignore real attacks – yielding “false negatives” – an obviously unacceptable approach.

False positives are not the result of poor software design by ID vendors. As Stefan Axelsson<sup>1</sup> demonstrated in his 1999 ACM presentation, there are some fundamental mathematical constraints that make false positives endemic to the whole paradigm of real-time signature recognition. Deviations from baseline norms can be caused by a variety of factors, many of them innocuous. **So false positives are inherently part of signature-oriented intrusion detection schemes or any other type of anomaly detection system.**

The unavoidability of false positives means that ID systems cannot be used to trigger automated corrective actions. After all, no one would want a system to potentially trigger the automatic shut down of access by an important customer who happened to be executing a lot of transactions that day.

**A second problem with ID systems is also related to their dependency on attack traffic signatures.** Whatever else may be said about attackers, they are a creative and constantly innovative breed. **Any system that relies exclusively on documented attack profiles will be vulnerable to new, as-yet-undocumented attacks.** While ID vendors are usually quite diligent about updating their products with new signatures, infosec managers may not want to risk having a window of vulnerability between the emergence of a new attack and the subsequent development, distribution and installation of a new signature.

**A third problem is that ID systems are fundamentally reactive.** When a real attack does take place, ID systems only alert security managers that something is amiss. It is then up to the security team to take remedial action. If the attack takes place at 2:00 AM, a substantial amount of time may pass between the alert and the remediation – leaving lots of cleanup work to be done and potentially allowing the attacker to do irreversible damage in the meantime.

Finally, as is clear from all of the points made above, ID systems are extremely administration-intensive. Highly skilled security professionals must constantly tune the system, update signatures, analyze alerts to determine if they are real or false, and then respond with appropriate remedial action. This is a lot of work for an uncertain amount of protection.

---

<sup>1</sup> S. Axelsson. *The Base-Rate Fallacy and its Implications for the Difficulty of Intrusion Detection*. In Proceedings of the 6th ACM Conference on Computer and Communications Security, pages 1-7, Nov. 2-4, 1999.

## **Honeypots**

A third security technology being used by many organizations is the honeypot. Honeypots lure attackers by presenting a more visible and apparently vulnerable resource than the enterprise network itself. Honeypots are useful for detecting attacks, since they provide a single point for security professionals to monitor for evidence of anomalous activity. They are also useful for forensics, since they can be specifically designed to retain data pertaining to an attack.

However, honeypots are not especially effective at attack prevention. With today's automated tools, attackers don't have to focus on a limited number of targets. **They can attack the honeypot and the enterprise network – and anything else in sight.** In fact, if they are incorrectly configured, honeypots can actually make the enterprise more vulnerable to attack by virtue of being linked logically to it.

Thus, while firewalls, ID systems and honeypots have their place in the arsenal of corporate defense, they are neither pro-active enough nor labor-efficient enough to meet the needs of today's IT-dependent, resource-constrained organization. Security managers simply can't afford to rely exclusively on these three technologies to protect their network perimeters. They need a solution that's more effective at preventing – rather than reacting to – all types of attacks while putting less strain on the time and energy of overworked security technicians.

## **Know Thy Enemy: Scans and Attacks**

The key to improving network security is to better understand attacks. **Attacks don't just happen. They are preceded by a phase of information collection.** Potential attackers scan and probe the target network for potential vulnerabilities to determine which type of attack to attempt. This is known as "reconnaissance."

Reconnaissance is an integral and essential part of any attack. To launch successful attacks, attackers need information about the topology of the network, about accessible network services, about software versions, about valid user/password credentials, and about anything else that will help them succeed in their efforts. Without such information, it is virtually impossible to successfully attack a network.

Unlike attacks themselves, reconnaissance can only be performed in some very basic ways. These reconnaissance methods may change subtly over time, but they inevitably share some basic attributes. Typical recon techniques include:

TCP/UDP port scan. This method accounts for at least 70% of all recon activity. The attacker operates at the network layer, mapping open TCP or UDP ports on a network host or hosts. This is extremely valuable information, since it reveals any applications running on the host that are accessible from the network.

NetBIOS probes. NetBIOS probes interrogate an IP host for computer names, user names, shared resources (such as shared folders or printers), and so forth. Responses to such probes will disclose the fact that the probed IP host actually runs a NetBIOS layer, and will reveal the objects sought by the attacker.

SNMP probes. These probes capitalize on the Simple Network Management Protocol (SNMP), which is used almost universally for communication between networked devices and management consoles. SNMP carries information about the

nature, configuration, topology, and health of those devices. As a result, attackers can gain a plethora of valuable information about all types of network resources.

Other recon methods include HTTP-based probes, “finger” probes, DNS zone transfers, and SMTP-based interrogation. Altogether, there are about twenty basic recon categories – all of which are well understood.

Typically, attackers use a variety of recon techniques. With each successive recon, the attacker gains more detail about the network’s vulnerabilities: an unpatched service, a visible NetBIOS resource, an open FTP port. Even when recon doesn’t yield any data, the attacker learns something about the network – i.e. that a host is not easily accessible. This helps the attacker further refine the attack strategy.

A typical attack can thus be viewed as having three stages:

- **the recon activity performed by the attacker**
- **the return of recon information to the attacker**
- **the attack itself, which is launched based on that recon information**

Understanding this three-stage attack process is central to effective defense. **In fact, security managers can take advantage of inherent flaws in the attack process to actually thwart attacks before they reach the firewall or the ID system behind it.** Just as attackers exploit vulnerabilities in the network to mount attacks, security managers can exploit vulnerabilities in the attack process to protect themselves.

## **ForeScout ActiveScout: Neutralizing Attacks Before They Reach the Firewall**

When network attacks are viewed in terms of the three-stage process described above, a natural question presents itself: **Why not respond pro-actively to the initial recon, instead of waiting for the attack itself?** Rather than waiting for the actual launch of an attack (i.e. stage three of the attack process), security managers should be able to respond immediately to pre-attack recon activity to preemptively neutralize any incipient threat to the enterprise. With this type of approach, attacks could be “nipped in the bud” before they compromised critical corporate assets. The network would only have to be defended against a finite number of well-known recon techniques, rather than an unlimited range of unknown attacks. And the issue of false positives would be virtually eliminated.

**This is exactly the approach implemented in ForeScout Technologies’ unique, ActiveScout™ solution.**

The ActiveScout solution is implemented in a security device situated behind the gateway router and in front of the firewall. From this location, ActiveScout can monitor all traffic heading to the corporate network. ActiveScout can be configured non-intrusively – via a line “tap” or a switch spanning port – allowing it to monitor traffic without introducing any performance degradation. Putting ActiveScout at the very edge of the network also allows it to perform its key attack-neutralizing functions, which will be explained below.

The ActiveScout solution uses patented *ActiveResponse*™ technology to proactively respond to attackers’ recon activity and thereby neutralize attacks in a uniquely powerful way. It does so using a three-phase paradigm that reflects the stages of the attack process.

## ***Phase 1: Receptor***

Most of the time, the ActiveScout functions as a passive monitor. **It quietly listens to incoming network traffic, looking for any signs of network reconnaissance.** This monitoring is done with a very high level of sensitivity, so that even very slow scans will be detected. This can be done because, as will be seen, false positives are not an issue under the ActiveScout paradigm.

During this stage, ActiveScout also sees which network services and resources are visible to the outside world (i.e. can be seen outside the firewall). This is done in order to avoid any interference with production traffic during phases 2 and 3. Essentially, ActiveScout can see anything a potential attacker might be able to see – while it also watches for attackers.

## ***Phase 2: Deceptor***

When reconnaissance activity is detected, ActiveScout automatically shifts to its active mode. It first identifies the type of recon being used by the suspected attacker. **ActiveScout then responds to the recon with information similar to that which is being sought.** For example, if a NetBIOS scan for shared folders has been detected, ActiveScout will respond with information about a shared folder using the NetBIOS protocols.

**However, the information supplied by ActiveScout is purposely counterfeit.** It looks exactly like the type of data that would have been supplied by a real target, but it is actually “deceptor” data provided with the express purpose of misleading the attacker. In the case of the above NetBIOS scan, for example, the shared folder name returned to the attacker would actually be a deceptor.

The potential attacker will view the information as valid, and will make use of it in any subsequent attack.

This deceptor data is very different from that supplied by a honeypot. Honeypots are *real* resources that are *accurately* pinpointed by recon activity. The deceptor data provided by ActiveScout, in marked contrast, gives the attacker *false* data about resources that do *not* actually exist. Also, ActiveScout’s deceptor data can specifically mimic all types of resources that may be targeted for an attack – such as an IP service or a NetBIOS resource. Honeypots do not provide this level of mimicry.

**It is important to note that up to this point, no alarm has been triggered.**

Security professionals do not have to respond to any situation or try to interpret complex traffic data. The deceptor data has been automatically sent to the suspected attacker and recorded in the ActiveScout’s database. The network continues to operate without disruption.

In most cases, the deceptor phase will be the last one in the response cycle. While almost all attacks start with a scan, very few scans actually result in an attack. A typical site can be scanned hundreds or even thousands of times per day, while there might only be a dozen or fewer real attacks during the same time period – so there will be no need for Phase 3.

However, neither the security team nor the enterprise loses anything by responding to these scans. There is no waste of infosec managers’ time or unnecessary

bandwidth utilization. In fact, it doesn't even matter if ActiveScout responds with deceptor data to traffic that turns out not to even be a scan at all. The entire process is completely innocuous in terms of live business being transacted on the network.

### ***Phase 3: Interceptor***

Some recon will be followed by attacks. Attackers in these cases will use the information gathered from their scans to launch their attacks. That information, of course, will actually consist of deceptor data provided by ActiveScout.

**Because the attacker is using deceptor data, ActiveScout will immediately be able to identify the attack when it occurs.** This is because – rather than depending on an attack signature –ActiveScout simply recognizes its own deceptor. In other words, ActiveScout has planted a “mark” by which it can detect and intercept traffic coming from a source that previously performed suspicious reconnaissance.

This approach is not unlike the use of marked bills to catch criminals in a sting operation. When the perpetrator tries to pass the bills, it provides evidence of their wrongdoing. By the same token, incoming traffic using deceptor data can be assumed to be offensive. It can thus be acted upon immediately and automatically, regardless of whether or not it conforms to any type of known attack pattern.

This is the point at which the ActiveScout system can generate an alarm with a high degree of confidence that a real attack is indeed launched. Alerts can take the form of e-mail, an SNMP trap, a line in a log, a pager message and/or any other appropriate type of message.

Depending on the company's security policy, all traffic from the offending IP address can be blocked for a predefined period of time as well. This blocking can be done by ActiveScout and/or in an integrated manner with the firewall<sup>2</sup>.

It's worthwhile to note that an attack may take place days or weeks after the scanning activity. The attack may also come from a totally different IP address than the scan. **The effectiveness of the ActiveScout solution is unaffected by either this time delay or the use of a “moving source.”** As long as the attacker performed a scan beforehand – which they invariably do – ActiveScout will perform.

## **The ActiveScout Advantage**

ForeScout's ActiveScout solution delivers numerous substantive advantages to today's harried information security professionals. These advantages include:

### **Proactive neutralization outside the firewall**

ActiveScout intercepts attacks before they can even reach the firewall – thus delivering maximum protection against business disruption.

---

<sup>2</sup> Prior to blocking traffic, ActiveScout must implement anti-spoofing mechanisms to make sure the source being blocked is indeed the actual source of the attack. One solution is to take advantage of the TCP three-step handshake process. In modern IP stacks, the TCP ISN (Initial Sequence Number) is very effectively randomized, so receiving a proper Sequence Number in a TCP handshake provides a very high probability that the source IP address is not spoofed – since the sending host needed to see the ISN in order to send out the next SN.

**Effective protection against both known and unknown attacks**

At no point does ActiveScout require any awareness of the specific nature of the attack being launched. It is only aware that there *is* an attack taking place. Since no previous knowledge of the attack technique is required, ActiveScout is able to identify attacks that are completely unknown to the security community.

**Elimination of false positives**

Since ActiveScout provides deceptor data to potential attackers and accurately identifies them upon their return, it never generates false alarms. Because ActiveScout functions automatically, it eliminates any drain on staff energy and productivity.

**Minimal labor requirements**

ActiveScout doesn't require continual signature updates, extensive tuning, log analysis, or other manual configuration. Plus, in the event of an attack, it automates the activation of immediate, appropriate counter-measures. No other security technology offers more protection for less work.

**Minimal impact on network performance**

ActiveScout is minimally invasive to the network connection – allowing production traffic to move without any discernible performance degradation.

**Low processing overhead**

ActiveScout can ascertain whether or not packets traversing the network are executing a scan without having to examine them in great detail. It therefore requires far less processing power than typically associated with ID systems or other sophisticated network security technologies.

The ActiveScout solution represents a radical innovation in information security technology and practice. By providing automated, intelligent and pro-active response to potentially malicious network reconnaissance, ActiveScout completely re-writes the rules of network security – in the corporation's favor. It thus represents a significant advance in the protection of critical business assets from the increasingly diverse and frequent external threats faced by information security professionals.

## **About ForeScout Technologies**

ForeScout Technologies is committed to delivering significant security innovations to enable organizations to more effectively protect themselves from a growing volume and range of security threats. The company's ActiveScout solution has radically improved enterprise information security by pro-actively neutralizing attacks through the identification of and pro-active response to attacker's pre-attack reconnaissance activity.

ForeScout is headquartered in San Mateo, Calif. The company's research and development center is based in Tel Aviv, Israel. ForeScout is a privately held company and has raised \$14M in funding from seed investors and leading venture capital firms Accel Partners and Pitango Venture Capital.

For more information on ForeScout Technologies, visit [www.forescout.com](http://www.forescout.com).