Getting the Knack of NAC:
Understanding Network Access Control

A Mirage Networks Industry Report
January 2006

Table of Contents

Part 1

Part 2

PART 3

Part 4

Getting the Knack of NAC:
Understanding Network Access Control

Part 1: What is NAC… and why do I need it?

I. Introduction: Network Security

A. Where We Are Now

Today's technology environment is defined by mobility. It's a productivity enhancement few organizations can be without – but the gain in productivity is causing an explosion of network security concerns.

Consider the dramatic increase in the number and capabilities of mobile devices: according to Gartner, the dominant trend in computer buying has shifted to notebooks, which now make up 29% of computers sold in the US and 31% of those sold worldwide.[1]. And not only are laptops becoming the computer of choice for many corporate employees, more and more IP-enabled devices are coming into the mix - PDAs, mobile phones, and gaming systems, to name a few, each bringing new security vulnerabilities onto the network. Further enhancing productivity – and jeopardizing network security – is the ubiquity of access. Whether at home, in a hotel, at a Starbucks, or even on a park bench, users require and expect access to corporate networks at a data rate that enables full productivity. The widespread adoption of broadband and wireless networking has made mobile computing the standard, not the exception.

This has created great challenges for IT and security professionals. Controlling the devices accessing the network has become increasingly problematic as these devices move in and out of protected corporate networks, and as the line between office and personal computer blurs or even disappears. And now, it's easier than ever for unmanaged IP devices to make their way into corporate networks.

This technology shift has IT security professionals asking two questions:

How do I control the access to my corporate networking resources?

-and-

How do I ensure that the resources that are allowed on my network aren't creating a security risk?

Before we can answer these questions, we must understand the roots of IT security.

B. How We Got Here

The correlation between productivity-enhancing technology and security technology is not new. As new technologies are adopted, criminal elements find ways to misuse them. Let's examine the origins of three prevalent security technologies: antivirus software, firewalls, and Virtual Private Networks (VPNs), the development of each driven by key advancements in non-security technologies.

1. Antivirus

In the early to mid 1980s, antivirus technology development was driven by the success of MS-DOS, and its impact on businesses and home PCs. A key new capability provided by personal computing was the ability to easily share and transfer files via floppy disks. As file sharing became standard behavior, the first viruses evolved to exploit it. This malware attached itself either to individual files or to the boot sectors of PCs to infect all floppy disks subsequently used on that PC. Antivirus technology arose and was widely adopted to preserve the value of data transfer through external media.

---

[1] ChannelRegister.com, *Notebooks keep becalmed PC market going*, August 2005

2. Firewalls

Firewalls are often associated with the advent of the Internet, but they actually came about as a result of networking and routing technology. As businesses began to connect their small departmental networks to larger shared networks, concerns arose about the ability of individuals to access computing resources and data on networks that didn't belong to them; the development of the firewall was the result. The firewall inserts itself as a barrier between a local trusted network and one or more external networks, regulating traffic between networks to prevent access to network and system resources from unknown or unauthorized sources. Connection to the Internet and its millions of worldwide users has made firewalls mandatory, and a standard part of virtually all networked environments.

3. Virtual Private Networks

VPNs, while not quite as prevalent as antivirus and firewall technologies, are found in almost all medium to large organizations. The need for VPNs was driven by two factors. First, corporations were looking for alternatives to expensive private networks connecting remote sites. Second, companies needed to enable their employees to connect to their corporate networks remotely. And of course, confidentiality of the data in transmission was critical: because this connection was going over open networks, it was susceptible to eavesdropping for both passwords and data. VPNs provided a mechanism to protect the confidentiality of the data and assure that the connection being made was legitimate.

These advances made the network perimeter stronger, acting as a moat between data and threats. But as perimeter security evolved, so did the methods designed to get at that data.

II. The Need for Something New

Today's threats are not entering corporate networks through the perimeter, protected as it is by the technologies reviewed above, and others, like intrusion detection systems. Rather, they are taking aim at the network's soft underbelly, through authorized endpoints, which, as known devices, completely bypass perimeter defenses.

A. An Example: the Zotob Worm

A recent example of this occurred in August 2005, during which the Zotob worm took advantage of a Microsoft Windows Plug and Play vulnerability. The worm infected PCs and propagated across networks by looking across random Class B addresses and sending a SYN packet (connection request) to port 445 on remote systems it found active. Upon finding a vulnerable machine, the worm exploited the vulnerability by downloading a copy of itself, infecting the PC and looking for other targets. Networks were flooded with traffic and crashed, costing organizations untold amounts in productivity.

Incredibly, this worm should never have been able to spread. Its propagation methodology required access to a TCP port through which other worms, most notably Sasser and Nimda, already had spread threats. Thanks to the notoriety of these other worms, almost every organization with an Internet connection blocked traffic to port 445, and assumed that they did not need to be concerned about Zotob. Despite this, Zotob quickly spread around the world, becoming one of the fastest spreading worms in history[2].

Beyond simply proving the cliché that assumptions are dangerous, this incident highlighted a crucial flaw in traditional network security: the ability, or rather, the lack thereof, to manage every device that plugs into the network. In the case of Zotob, for example, the worm entered organizations when known, infected mobile PCs entered the network, either directly or through VPN connections.

---

[2] Slashdot.com, *Zotob Worm Hits CNN and Goes Global*, August 16, 2005

Which brings us full circle to the question of controlling endpoints to minimize security risks without impeding business.

III. Voilá: NAC

Network Access Control (NAC) aims to do exactly what the name implies: control access to the network. It is still an emerging technology space, and many vendors are taking advantage of this lack of definition to jump on the NAC bandwagon. But if we boil down NAC to its essence, we are referring to the ability to:

- Enforce security policy and restrict prohibited traffic types
- Identify and contain users that break rules or are noncompliant with policy
- Stop and mitigate day zero and other threats

In the beginning…
Early on, Cisco and Microsoft created their own proprietary standards that led to the creation of key NAC concepts, prompting standards bodies to influence the direction of further development.

A. Proprietary standards

The first standards developed were those of vendors, recognizing the need to test devices that enter a network to see if they meet a baseline security policy.

1. Cisco NAC

Cisco initiated its Network ***Admission*** Control (NAC) program in 2004 in an effort to provide a mechanism to check endpoints as they enter the network. As defined by their initial program, there are three key elements of Cisco's NAC.

1. The first element is the reliance on endpoint software to provide both security functionality to the endpoint and an integration point to the NAC infrastructure. This requires two endpoint agents: the Cisco Trust Agent (CTA) and Cisco Security Agent (CSA). The CTA communicates with Cisco NAC-enabled security agents to gather key security information and pass it on to the Cisco infrastructure. The CSA is Cisco's recommend security software, although antivirus software and/or personal firewalls can be substituted.

2. The second element of Cisco's NAC strategy is contained in the software of their Network Access Device (NAD), i.e., a router or a switch. After upgrading, the Cisco networking operating system is able to recognize endpoints on entry to the network and check them for the presence of the CTA. Upon finding that agent, it gathers security information relevant to the state of the device entering the network. The NAD can then forward this information to the third element of their NAC strategy, the policy servers.

3. The third element, the Cisco policy server, consists of their Access Control Server (ACS) (which is a RADIUS Authentication, Authorization, and Accounting (AAA) server) that can grant access to the network based upon the results of the host agent results. In Cisco's homogeneous solution with CSA, the information passed to the ACS server is relevant to the Operating System version and patch level. Through integrations with antivirus vendors and other endpoint security solutions, it can also check such things as AV signature level, if security agents had recently found or mitigated threats, and other security policy information. Once a system has passed these checks, it is allowed on the network with an appropriate level of access. Different access levels are usually enforced by assignment to various VLANs.

Cisco rolled out its NAC solution initially on Layer 3 devices (routers) for protection and checking of remote access devices. It has since launched a solution that supports Layer 2 capabilities (switches). As with all Cisco initiatives, for Cisco NAC to function optimally, it must be in a homogenous Cisco network, with new, NAC-ready infrastructure elements.

2. Microsoft NAP

Microsoft was the second major vendor to deliver its own NAC standard. Called Network Access Protection (NAP), it was designed to check the security policy compliance of Windows devices entering the network. There are three distinct aspects of NAP: health policy validation, health policy compliance and network isolation.

1. For health policy validation, an endpoint must have Microsoft's or a supporting vendor's System Health Agent (SHA), which checks the current security state of the device, such as the status of OS patches and antivirus updates. The SHA reports to a System Health Validator (SHV) that determines whether or not the endpoint is adhering to corporate security policy.

2. After policy validation, health policy compliance features may be used to fix any identified shortcomings of the device. In this case, after the system is identified to be non-compliant, it can be fixed automatically using services such as Microsoft's Systems Management Server (SMS).

3. Finally, network isolation can be implemented through one of several enforcement options. Currently these options include DHCP quarantine through the use of a DHCP Quarantine Enforcement Server (QES), VPN quarantine using a VPN QES, 802.1x quarantine where the network access device can limit the access of the endpoint to specific protocols or VLANs, and finally IPsec quarantine using an IPsec QES. Each of these isolation options have certain capabilities and limitations that can be explored further on Microsoft's website.

Microsoft NAP is expected to be delivered and available in conjunction with the release of Microsoft Vista in late 2006. As with all Microsoft initiatives, NAP will work optimally only for networks and devices running on a Microsoft OS.

B. Open standards

In response to Cisco's and Microsoft's approaches, the Trusted Computing Group (TCG), through its Trusted Network Connect (TNC) subgroup, kicked off an industry initiative to influence the development of vendor-neutral NAC solutions.

1. TCG TNC

The goal of the TNC was to establish a standard for ensuring host integrity of devices connecting to the network, where integrity indicates that the device is both free from malicious code and that the device is up to date in its own protections against threats, with current patches and the like. The TNC architecture can be broken into three major components; the Access Requestor (AR), which is the device entering the network; the Policy Decision Point (PDP), which compares the AR's credentials against an established security policy; and the Policy Enforcement Point (PEP), which can grant, deny, or limit access for the AR.

The configuration of the AR includes a TNC Client (TNCC) that runs on the AR and interfaces with another software application, the Integrity Measurement Collector (IMC) one or more of which runs on the AR. The IMC checks for antivirus signature level, operating system patches and other host based security measurements. The TNCC aggregates inputs from the IMC(s) and communicates this information to the PDP. The PDP uses this information to analyze the host's security measurements determine their compliance to established security policies. Once the AR's adherence to or variance from policies is determined, the PDP decides what level of network access the AR should have. The PEP is used to enforce these decisions. A PEP almost always resides in the network and is part of the entry process in some way: PEP enforcement mechanisms can include integrations with the Authenticator in 802.1x, with AAA servers during the authentication process, and with secure DHCP servers for the assignment of IPs and VLANs. They can also be implemented in the network infrastructure itself, either through a network security appliance or through integration with the network access device and the use of Access Control Lists (ACLs).

C. Independent Vendor Solutions

As standards-based approaches to NAC become accepted, many vendors who play in adjacent markets are taking the opportunity to claim their spot, integrating into the above frameworks and performing part of the NAC process. Most of these solutions take the role of either or both a PDP and/or a PEP as described by the TCG TNC architecture.

Many of these solutions have gone beyond the initial concept proposed by Cisco to provide a more complete approach to NAC throughout a given endpoint's access lifecycle, an approach that analysts like Gartner[3], advocate.

IV. Summary and Next Steps

Mobile technology is rapidly rendering the network perimeter moot; protecting the "soft, chewy center" is critical. Designed for exactly this environment, the NAC approach is a promising one.

There are myriad NAC solutions on the market. Deciding which one works for a given organization can be challenging. Key elements to look for include pre-admission NAC, post-admission NAC, and quarantine and remediation capabilities.

In Part 2, we will examine pre-admission NAC in detail; in Parts 3 and 4, we will focus on post-admission NAC and quarantine and remediation and key considerations you should keep in mind when deciding on what NAC strategy will work for your organization.

[3] Gartner, *Protect Your Resources With a Network Access Control Process*, December 2004

Getting the Knack of NAC:
Understanding Network Access Control

Part 2: Pre-Admission NAC

I. Introduction: Pre-Admission NAC

In Part 1 of this Paper, we examined what has driven the development of security technology, and how that has created the situation in which we find ourselves today. The goal of this report is to answer the two key questions IT security professionals find themselves asking with increasing frequency and urgency:

How do I control the access to my corporate networking resources?

-and-

How do I ensure that the resources that are allowed on my network aren't creating a security risk?

We also introduced the three core elements of a successful NAC implementation: pre-admission, post-admission, and quarantine and remediation. In Part 2, we focus on the first of these elements.

II. Pre-admission NAC

As you will recall from Part 1, the three main NAC standards (Cisco NAC, Microsoft NAP, TCG TNC) concern themselves with pre-admission (a.k.a. on-entry) NAC. Pre-admission checks are critical, but assuming that clean devices cannot become infected or hacked once ON the network could well be considered myopic. Post admission infection is addressed later in this document.

For a NAC solution to be effective, it must deliver two essential pre-admission capabilities. First, it must be able to identify a new device connecting to the network. Second, it must be able to test the endpoint for adherence to security policy and restrict access for those devices that do not meet defined entry criteria. Together, these capabilities should provide data that can be used to compare a device's current security state against established security policy criteria, to determine how much or how little access that device is allowed.

A. Endpoint Detection

1. Endpoint detection: software approaches

There is a significant variety the approach NAC solutions take to both identify and control access to the network.

a. 802.1x

A common approach to identifying **and** controlling network access for a user on a device is through the 802.1x standard for network authentication.
There are three elements in this process: the supplicant, the authenticator, and the AAA server.

This standard defines a communications protocol between three elements: a supplicant, an authenticator, and an authentication, authorization, and accounting (AAA) server. The supplicant is software deployed on a user's workstation, **and** the authenticator is implemented on a network device such as a router, switch or wireless access point. The authenticator requests information from the supplicant and passes that information to the AAA server.

802.1x uses the Extensible Authentication Protocol (EAP). EAP is used to communicate authentication information between supplicant and AAA server through the authenticator, and typically uses it to send additional host integrity data to the authenticator.

i. 802.1x: pros and cons

On the plus side, 802.1x detects entry to the network before the endpoint has an IP address **and** can cause havoc on the network.

The drawbacks are tied to manageability. To be effective, the 802.1x approach requires both management of every endpoint device and integration with a third party: all network devices must be configured for 802.1x, and they must also be integrated with the authentication server. Additionally, this approach often requires an upgrade of legacy networking equipment and OSs. When you consider the reality of what these management requirements will demand in terms of time, in addition to the cost and time required to upgrade the infrastructure, it is clear why adoption of the 802.1x has been slow.

b. Dynamic Host Configuration Protocol (DHCP)

Another pre-admission approach for identifying endpoints **and** restricting their network access is through integration with a DHCP server, very commonly a DHCP Quarantine Enforcement Server (QES). In this approach, the device accesses the network **and** sends out a DHCP request for IP address assignment. DHCP solutions usually require a host agent – software on the device itself – that serves as what is known as the Policy Decision Point (PDP). In NAC solutions using this approach, there is usually also an in-line DHCP appliance that deploys between the organization's DHCP server and the network.

When a device requests an IP address, the DHCP NAC appliance operates on the "better safe than sorry" assumption and assigns an address that sends the requestor to a restricted, quarantine VLAN. If the device has the DHCP agent, the DHCP server queries it to ensure that the agent is in compliance with policy. If it is, the device is assigned a replacement IP address, giving it access to the appropriate network VLAN. However, should the device either have no agent, or an out-of-compliance agent, it can either be blocked from entry entirely, be kept in a quarantine VLAN, or be required to download an agent for testing of unmanaged devices.

i. DHCP: pros and cons

DHCP is often a good choice for NAC because for the most part, the majority of valid company users use DHCP for IP address assignment and therefore is a good hook for recognizing new devices. Another advantage of this approach is that it works regardless of whether or not local network devices support the 802.1x standard or not and it is generally easier to deploy than the 802.1x implementation.

There are several drawbacks to this approach. As with 802.1x, requiring devices be equipped with agents, and that those agents be kept up-to-date, is problematic in and of itself. Further, most of the agents are restricted to Windows clients.

But in the shadows loom two greater dangers. First is the issue of static IPs, which easily bypass DHCP security. This could be a user assigning his/her device a static IP address – a tactic employed both by cyber criminals and valid users wanting to avoid the entry check – as well as devices with static IP, like servers or printers. Second, by using broad networks for quarantining high-risk devices, this approach creates a "VLAN of Death" – a virtual leper colony in which devices ultimately end up infecting **and** re-infecting each other.

c. Authentication

A third pre-admission approach to identifying **and** controlling network entry is through integration with the user authentication process; the approach that Microsoft has taken with their NAP initiative is an example.

There are two types of solutions taking this approach. In the first, a replacement or proxy for authentication is implemented in-line with the actual AAA server; in the second, the in-place authentication server is upgraded with software that can recognize both a host's authentication and integrity credentials (Microsoft NAP is an example of the latter). In either instance, endpoint integrity determines what network and system resources the endpoint may access.

Typically, these solutions have to integrate with a Policy Enforcement Point (PEP) for mitigation, using 802.1x or a secure DHCP QES or VPN QES for network restriction. These systems rely more on the server infrastructure than the networking infrastructure for triggering a host integrity check.

i. Authentication: pros and cons

This approach is useful for restricting access to network servers and services on the network. However, it does have one critical flaw that leads to a dangerous irony: unmanaged devices don't initiate authentication server identification, which means no checks will be run on those devices. By definition, unmanaged devices are higher-risk devices, so the devices that most need a check are the devices that won't get one.

2. Endpoint detection: appliance approaches

Network security appliances are an alternative to the above to detect the entry of the endpoint to the network **and** initiate policy checks. Their deployment is what differentiates them: in-line vs. out-of-band.

a. In-line appliances

In-line appliances either replace the edge switch or sit in line directly behind the access layer switch. This allows them to see traffic, detect devices **and** initiate policy checks. Alternately, they can be deployed to sit in-line with either the DHCP server or the Authentication server, and then function in a manner similar to that discussed in section 1b above. This similarity is the reason that the latter deployment model is less desirable, as it makes the security easy to bypass.

i. In-line appliances: pros and cons

The advantages of an in-line NAC appliance is that it will see all endpoints coming on the network, and can control their traffic effectively because all the traffic flows through the appliance. The drawback of this approach is that it is essentially replacing the access or distribution layer switches, which can be a costly, **and** time consuming process in terms of network upgrades **and** potential rearchitecture.

b. Out-of-band appliances

A final approach the use of a dedicated security appliance that is deployed out-of-band. To deliver optimal results, the appliance must be capable of maintaining a real-time IP address map **and** be able to detect when a device needs to be tested **and** controlled. Here, too, there are varying approaches. Some use an alternative PEP, such as the DHCP or Authentication servers, or 802.1x as described above.

The best take the "less is more" approach, eschewing third-party PEPs for the increased cost and latency they will deliver. The Mirage Networks NAC solution is an example of this approach. It uses a unique combination of maintaining an active IP map by monitoring traffic at both Layer 2 and Layer 3 of the network **and** using ARP management to isolate endpoints from the rest of the network for testing.

i. Out-of-band appliances: pros and cons

The positive aspect of this approach is that a device can be isolated from the network without the ability to bypass NAC checks, as is the case with DHCP and Authentication hooks, and without costly network replacement or reconfiguration that we see with 802.1x and in-line solutions.

B. Device Testing for Risk Profiling

Once a device has been detected, it can be tested. Testing can take many different forms, from simple authentication **and** services checks to deep scans of the registry **and** applications to determine if the device's operating system (OS) and security relevant applications are fully up to date and adhere to network security policy

It should be noted that entry scans are better at providing risk information than infection information. In other words, pre-admission NAC testing can determine if a device is out of compliance with security policy or is vulnerable to be threats, but won't provide a great deal of information on an actual infection. This is why post-admission monitoring– something we'll discuss in detail in Part 3 – is so critical.

C. Policy adherence

Most NAC solutions have some element of testing endpoints for policy adherence as a requirement for admission to the network. Policy testing identifies endpoints that fail to meet some criteria established by the IT or security administrator. These tests can include checks for services or applications, i.e. peer-to-peer (P2P) software or Web servers, OS patch levels, AV signature levels, **and** processes to identify unmanaged, and/or unauthenticated devices or end users.

1. Active services/applications

An elemental endpoint check is looking at running processes **and** available services that may either be restricted, like P2P applications, or required, like AV software **and** other security applications. Some of these checks can be run external to the device from the network, **and** will look at the network ports **and** services that are available on the endpoint. This technique can be used to identify Web, FTP, and/or P2P servers, **and** many other applications that have a networking component. The presence of some services can also indicate the absence of certain security software, like a personal firewall, for example. If an endpoint has networking services available for connection – like TCP port 21 for FTP connections – that are restricted by a correctly configured corporate personal firewall, then one can assume that the personal firewall is either not running or is improperly configured.

If checking for services that have no available networking component is important, some application-level checks require software deployed on the endpoints for accurate endpoint examination. In this case, the endpoint security software must interact with a network-based policy server that serves as the PDP to ensure that the results reported by the endpoint agent match the policy criteria established by the policy server.

2. OS patch-level vs. vulnerability checks

There are two approaches to OS patch levels checks **and** the subsequent vulnerability state assessment of an endpoint. In the first, a minimum patch level is tested on endpoint devices entering the network. Those that don't meet this minimum patch level are either sent to quarantine or are denied access to the network. The advantages of this approach: it is a very objective measure of the requirements for entry to the network, **and** it can eliminate devices that may be vulnerable to a certain type of attack. There are drawbacks, however. First, in network environments that enable minimum IT oversight **and** allow the use of a wide variety of devices, it can often be difficult, **and** sometimes virtually impossible, to require the very latest patch level for every endpoint that enters the network. If highly restrictive policies are enforced, they can often result in a high volume of Help Desk calls **and** an unhappy user population. Second, to really check for OS patch levels without opening up endpoints to dangerous queries, this approach requires endpoint agents that come with their own management overhead **and** headaches.

A second approach is to test the endpoint for vulnerability to attacks from the network. This approach examines what OS patching is intended to prevent, i.e., vulnerable services that can be exploited by threats. An advantage of this approach is that an endpoint can be checked from the network, without agents, for vulnerabilities to self-propagating threats. A drawback of this approach is that it doesn't identify systems that need patching during the entry process. It is possible for a system to be protected against a threat, but still not have the OS patch fixing a particular vulnerability. This could be the case if a personal firewall were in place, for example. In many instances, the use of one or both of these techniques is a philosophical decision based on risk tolerance **and** the amount of management overhead the administrator is willing to take to use one or the other technique.

3. Antivirus signature level

Another common check by NAC products is for the antivirus (AV) signature level. The concept behind this approach is that devices that are up to date on AV signatures are unlikely to be infected with a threat. An endpoint agent interacting with a network policy server enables these checks. The advantage of this type of check is, obviously, that up-to-date AV software is good protection against threats. But there are several "gotchas" here, if this is the only policy check for admission being performed (as is often the case with many Cisco NAC deployments, in which the endpoint agent is AV software from a major AV vendor.) First is defining the minimum signature level for entry: if this is defined as the most recent, then systems will be routinely blocked during entry as they are required to update their AV software. This can lead to a very unsatisfactory user experience. Second, even with updated AV, devices may still be infected with a threat that will go undetected until post-admission, as a full system scan usually takes tens of minutes **and** is an unacceptable delay for entry to the network. Finally, AV signatures, by their nature, are always released <u>after</u> an outbreak. This means that there is a period of time that the network is vulnerable to attacks from the latest threat. The upshot: AV signature checks play an important role in admission criteria, but usually aren't sufficient to ensure that infected devices don't make it into the network.

4. Managed vs. unmanaged devices

A final policy check for endpoints entering the network is determining its status as a managed **and** authenticated device, as opposed to one that has not been seen before or approved to be on the network. This check is really all about controlling access to the network for machines that are not company-approved or managed by the IT staff. This approach can be used to identify **and** test contractor machines, home machines, **and** many handheld devices that may make it onto a corporate wireless network. Many times, this is implemented to ensure that any new machine that enters the network has undergone certain checks, which may take longer **and** be more rigorous than standard checks, before they are allowed on the network. In this case, they may be sent to a quarantine area of the network, where these tests **and** registration processes can be performed before network admission is allowed.

B. Threat profile

Another type of test that can be run on network admission is threat checking. Unlike policy checks, which give an indication of the risk profile of an endpoint, these tests look for actual threats.

There are two basic approaches for testing for endpoint-originating threats. The endpoint can either be scanned for threats on entry to the network, or it can be monitored from the network to identify threats coming from the endpoint.

1. Infection testing

Testing for threats on endpoints usually involves a network-based AV scan that occurs in the context of a Web session. These tests are usually pared down to minimize user latency, **and** look for the most recent **and** most common threats. This approach can be helpful in identifying **and** removing actual threats on endpoints on entry to the network. The reason this approach is generally not used, especially in the LAN, is that the less time the scan takes the less effective it is, **and** the more time it takes the more unacceptable it is as a component of the admission process to end users. This approach also suffers from the fact that AV signatures are developed **and** deployed after a new threat has already reared its head, **and** so cannot catch early, or 'day zero,'" threats.

2. Threat detection

A second approach to threat profiling is examining traffic originating from an endpoint, to identify and stop threat entering the network from the endpoint itself. This can include both signature-based and/or behavioral-based detection through analysis of the traffic originating from the endpoint. The advantage of this approach is that it is transparent to the end user **and** usually monitors the traffic from the endpoint both on admission, as well as post-admission to the network. The drawback of this approach is that, just like AV solutions, signature-based solutions need to be updated frequently, **and** there is a risk of false positives on many systems that monitor behavioral anomalies. It is also important to look at the speed of detection, **and** the speed of mitigation on detection of a threat, since many threats spread in a matter of seconds.

III. Summary and Next Steps

The approaches to pre-admission NAC are almost as varied as the threats that those solutions are designed to nullify. IT security pros must examine carefully which “pros” they are willing to trade for which “cons” to find the solution ideal for their networks.

Of course, pre-admission is only one piece of the puzzle. For a robust NAC solution, post-admission NAC, and quarantine and remediation capabilities are also necessary.

In Parts 3 and 4, we will focus on these remaining two elements of a successful NAC implementation, and key considerations you should keep in mind when deciding on what NAC strategy will work for your organization.

Appendix: Comparison Table: Endpoint Detection Methodologies

Detection and Control Method	**Pros**	**Cons**
802.1x (Layer 2)	• Authentication of device upon network connection • Interacts with switch for port level mitigation **and** VLAN assignment	• Not all switches support 802.1x • Configuring 802.1x across an enterprise is time consuming **and** difficult to support • Not all devices authenticate (printers)
DHCP (Layer 3)	• Easy to intercept requests for IP assignment • Less devices are required because it only needs to be in-line with the DHCP server	• Control is easily bypassed by devices assigning a static IP • Won't catch devices behind a NAT device (i.e. routing WAP)
Authentication Integration	• Can couple authorization to network resources, i.e. servers, printers, etc., to the authentication process **and** results of host integrity check. • Can be both identity **and** device specific, i.e. stolen laptops aren't a problem if they can't authenticate	• Devices often don't have to authenticate to have access to network **and** spread threats • Not a strong solution for unmanaged PCs
Network Security Appliance In-Line	• Much like 802.1x these devices see a device as soon as they are on the network • If deployed behind or as a replacement of the Access layer switch these devices provide strong control of network traffic traveling through them	• This approach can be extremely costly **and** time consuming. • Ideal deployment is behind every Access layer switch **and** requires network re-architecture • Adds a point of failure on the network
Network Security Appliance Out-of-band	• No network re-architecture required • Appliances like Mirage's CounterPoint can detect every host entering the network **and** test them without bypass capabilities seen in DHCP **and** Authentication approaches	• Some devices don't have their inherent mitigation devices **and** must use one of the above technologies • Some devices that only watch IP traffic **and** don't see Layer 2 traffic may be late in detecting **and** assessing a device.

	Cost to Implement	Difficulty to Manager	Vendor Support Problems	Vulnerability to Attack
802.1x	Med	Med	Med	Low

Getting the Knack of NAC:
Understanding Network Access Control

Part 3: Post-Admission NAC

I. Post-admission NAC: Introduction

As previously reviewed, Network Access Control (NAC) is a process that encompasses several phases in controlling devices' access to network resources.

In Part 2 of this report, we discussed various methods to retain control over which devices are allowed entry access to the network. In Part 3, we will take that to the next logical step, and discuss how to control which devices are allowed to retain access while on the network, by monitoring endpoints for policy violations and risk of infections. Part 4 will delve into the specifics of quarantine and remediation of at-risk and problem devices, whether that is on entry during risk assessment or post entry during threat mitigation and policy enforcement.

II. Decision 1: In-Line vs. Out-of-Band

The first decision to make in selecting technology with post-admission security capabilities is whether an in-line or an out-of-band solution will serve you best.

A. In-Line Solutions

In-line solutions have the advantage of seeing all traffic passing through a flow point and being able to surgically drop bad traffic and isolate traffic from an offending endpoint. This is a significant improvement over historical Intrusion Detection Systems (IDSs) that had limited mitigation options once a threat was detected.

There are some drawbacks, however. Because they are a part of the network infrastructure and can constitute another point of failure on the network, high-availability capabilities are critical and evaluations of this technology usually include network latency tests under load and threat outbreak. A second consideration for in-line devices is that their mitigation options are limited to traffic flowing through the device. Here, their strength can often become a key weakness, especially when that device is deployed at the core of the network. In the instance of an outbreak, they can protect bad traffic from flowing out of a given segment, but threats can often have significant impact on internal network segments where the infection may spread from one device to a peer on its own broadcast network, never passing through the in-line device. In this manner they can propagate on their broadcast segment and to adjacent VLANs that are not protected by the in-line device, causing significant damage.

B. Out-of-Band Solutions

On the flip side, we have out-of-band solutions, which typically sit off a mirror port on a switch and monitor network traffic for threats. Compared to their in-line alternative, out-of-band solutions historically have posed a far smaller problem in terms of adding points of failure.

Where their drawbacks have typically come is in the form of threat mitigation. Most out-of-band solutions must rely on using the networking infrastructure to mitigate threats. This is accomplished by using either a networking standard like 802.1x, or by maintaining an active session into networking equipment for the purposes of changing Access Control Lists (ACLs) for mitigation. Unfortunately, many networking devices in deployment today are not configured for 802.1x, and often they simply do not support the latest mitigation protocols. Additionally, when ACLs are used extensively, there, is the added concern of latency, as well as security concerns when credentials are required for the entire networking infrastructure for mitigation.

III. Decision 2: Agent-Based vs. Network-Based

Unlike pre-admission NAC, which is simply a risk assessment that measures the potential risk of a given endpoint that is connecting to the network, post-admission NAC detects real threats and active policy violations, so the at-risk devices can be identified and fixed.

A. Your Choices: Agent-Based or Network-Based

A key element of post-admission NAC is the ability to recognize and stop undesirable activity – threat propagation or security violations – originating on a given endpoint. Two primary approaches for accomplishing this objective the objective of monitoring and responding to endpoint activity have emerged:

- agent-based: deploying security technology that resides on the endpoint (host)
- network-based: implementing security on the network

Broadly stated, agent-based approaches protect the individual endpoint from threats on the network, and network-based approaches protect the network and network-attached assets from the threats on individual endpoint.

An obvious and critical ramification of the agent-based approach is that, to provide sufficient post-admission security from endpoint-originating threats for the network as a whole, the device population must be 100% managed and in near or total compliance with the latest version and configuration of the host-based software, which still does not address day-zero attacks. The network-based approach, on the other hand, works well at defending the network regardless of the state of an individual host, and can generally protect the network in the instance of an outbreak that occurs despite best efforts to secure the endpoint. Network-based security cannot, however, protect the host when it is outside the protected network environment. As a result, many organizations choose to deploy both technologies to achieve "best of both worlds" security.

1. The Agent-Based Approach

The upshot of the above is that the agent-based NAC approach assumes that the network will be protected post-admission because all endpoints are running their own protective software to prevent infections. To make this assumption a reality, several requirements must be enforced.

a. Disallows Network Access to Out-of-Policy Endpoints

The first requirement for post-admission NAC to be effective is that the pre-admission NAC solution enforces a policy that vulnerable endpoints, that is, those that are not complying with policy and are therefore at risk, are not access to network resources in the first place.

Certain considerations for this criterion make it a difficult one to guarantee. First, there must be an agent on every IP-enabled and network-attached device. When the modern enterprise includes IP devices like printers, phones, PDAs, and laptops, and operating systems such as Windows, Linux/Unix, Solaris and Mac (and that's just the tip of the diversification iceberg!), getting 100% coverage of devices with agents is at best challenging, and usually, not possible.

Second, administrators must have a strategy that can detect and stop endpoints on which the protective software is stopped or reconfigured, which usually occurs for one of two reasons:

- a worm or other malicious code that makes its way onto the computer may disable the protective software (many worms target AV and personal firewall software in this manner)
- users who have administrative rights to their own computers often make changes to security software that slow or otherwise interfere with their daily operation of that device

b. Checks Security Configuration Of Host Prior To Admission

Another consideration for agent-based security is that it must also confirm not just that the security software is running and up to date, but also that the security software is configured in accordance with a standard desktop image and policy configuration.

2. Agent-Based Security Technologies At A Glance

Once a company has decided on an agent-based security approach, it has several different technologies from which to choose At this time, the savvy IT manager will determine which IP devices can be protected by the vendors and technologies in play.

a. HIPS

Host Intrusion Prevention Systems (HIPS) are typically security software that install deeply within the host, wrapped around the OS kernel, where it serves as an intermediary for system calls and other operations. HIPS solutions can be a powerful tool in preventing infections on a given device, and are particularly effective at stopping categories of attacks, such as buffer overflow attacks, in which a new and unpatched vulnerability is exploited. HIPS can also regulate which processes are allowed to run on a given system, and which fall outside of that “white list.” They can also enforce and restrict access to other system resources, such as IP services like file sharing and Web services.

The biggest drawbacks of HIPS are the ongoing administration of the software and the user experience. HIPS generally require a significant amount of ongoing care and feeding for the security policies that are being implemented. The more locked down the user environment, the more time that is usually needed by the IT administrator in making exceptions to the configurations for individual users. The more open the environment is when deploying HIPS, the more burden is shouldered by the individual user with respect to making security decisions on an ongoing basis. This can also result in increased help desk calls and dissatisfied users.

HIPS are particularly suited where a highly locked down environment is enforced and where individual users have very little, if any, say over the configuration of their individual PCs.

b. Personal Firewalls

Personal firewalls started out primarily as a mechanism for restricting access to system TCP/IP and other network services on a given host. Today, they have evolved to also perform duties of a policy management system that can check system configurations of other security software, like AV, and can check the status of the OS patch levels and other security configurations. Personal firewalls typically don’t get as deep into an endpoint’s OS as HIPS, but the line between these two technologies is becoming blurred into a broader endpoint security product set. Many of the advantages and disadvantages described above with respect to HIPS solutions also apply here, especially with respect to the IT administration and the user experience.

c. Patch and Policy Management

A final category of security software used for post-admission NAC is patch and policy management software. Unlike HIPS and personal firewall solutions, this software doesn’t aim to provide an active defense against threats it may not have seen before. Instead, it relies on OS patching and having the appropriate security configurations and software running to secure the machine. As such, patch and policy management systems don’t provide much in the way of post-admission defense, although they can and do provide significant value in pre-admission NAC, and serve a valuable function in ensuring that an endpoint is up to date and individually in accordance with a defined security policy.

3. Network-Based Approach

An alternative to agent-based security software is to deploy security technology in the network to mitigate threats and enforce policy post-admission. Like agent-based technologies, this approach also has specific characteristics which should be noted.

a. Protects The Network From Potentially Infected Hosts

The objective of network-based solutions is to identify and stop threats and policy violations originating from an endpoint. The assumption here is that it is inevitable that an infected or vulnerable device will at some point enter the network. In this scenario, the network security solution must quickly and effectively be able to identify even ”day zero” threats – those which have never been seen before.

b. Does Not Require 100% Compliance to Protect Network

Network-based security solutions do not assume 100% of network-attached devices are in compliance with security policies. This is especially important in environments where there are unmanaged devices on the network – home PCs, IP phones, contractor laptops, or student PCs, and the like – where not all IT assets with access to the network are company owned. Additionally, there are the “headless” endpoints, like networked printers and fax machines, on which agents cannot be deployed.

The drawback to a pure network-based approach is that the endpoints themselves are not protected when they are off the corporate network, or from risky user practices on the network.

c. Monitors the Network Continually for Threat Outbreaks

Network-based security solutions used for post-admission NAC are typically a variant of Intrusion Prevention System (IPS), behavioral security, or anomaly detection systems. All of these systems offer the capability to detect threats. Most modern network-based solutions that are appropriate for deployment in the network interior also have an enforcement capability that allows them to isolate and quarantine a single offending endpoint when a threat is discovered, or that at least can drop and block malicious traffic on the network.

Approaches differ, with the behavioral branch of the network-based security tree offering more variety than a signature/vulnerability-based approach. The quickest way to determine the true nature of a security solution is by examining whether it relies on updates to its system in order to catch and stop the latest threat or to protect against the latest vulnerability. If regular updating is required, that system typically cannot make claims to being behavioral-based or to delivering day-zero protection.

Of the behavioral solutions, there is a distinction between anomaly-based solutions and behavioral threat detection solutions. Anomaly-based solutions rely on profiling the network activity over a period of time and then detecting “anomalies” in endpoint behavior that deviate from the established normal baseline. These systems have typically evolved from perimeter solutions focused on Denial-of-Service attacks into interior threat protection systems. One consideration of these types of systems is the effectiveness in a user environment that may have a diverse profile of activity. Anomaly-based solutions work best in environments where there is little variance in services and daily activities from one host to another. In organizations where there is a great deal of variance, these systems must be tuned down to account for the normal variance in activity, rendering their detection of threats less effective. If the system is tuned to threat detection first, then the IT administrator will often deal with large volumes of false-positives.

The other type of behavioral solution deals specifically with detecting behavior that is indicative of threat propagation. These types of systems generally do not need to know a specific threat in order to stop it.

Both this and anomaly-based detection have some drawbacks stemming from their approach. First, they typically can’t, or aren’t, very good at, naming individual threats they find on the network, since they don’t use signatures: the theory here is that a threat by any other name would still be as dangerous, and that stopping its propagation is more important than naming it. Second, they can often have problems at detecting an insider threat in which a user exploits a specific vulnerability in a known network system resource.

IV. Summary and Next Steps

We've identified the different categories of post-admission NAC technologies, as well as the drawbacks and advantages of each. It behooves IT administrators to take a close look at any NAC solution that promises 100% effectiveness: there is no such thing.

There are, of course, some products that are better than others. These are usually dependent on the innovation of vendors who recognize the pros and cons of available technology and seek to "build a better mousetrap."

Mirage Networks, for example, delivers an out-of-band solution that does not have the usual drawbacks inherent in the out-of-band approach – it does not require 802.1x or changing ACLs to detect and mitigate, and uses a unique approach to detect and stop attackers by way of ARP management, a concept which is discussed more fully in Part 4.

Getting the Knack of NAC:
Understanding Network Access Control

Part 4: Quarantine and Remediation

I. Introduction: Quarantine and Remediation

The pre- and post-admission NAC capabilities we've discussed so far can be impressive, but without the ability to remove an offending endpoint from the network and fixing what ails it, it's akin to a burglar alarm that lets you know someone's in the house, but doesn't call the police. Informative, but odds are good that you won't like how the situation resolves itself. And that's where quarantine and remediation come in.

II.Quarantining Options

Quarantining options have evolved significantly in recent years, expanding on the less complex, "old-school" capabilities that simply reset a TCP connection or blocked an offending endpoint on the router with an ACL.

A. DHCP integration

Some solutions, and most notably the standard proposed by Microsoft NAP, can rely on integration with a DHCP Quarantine Enforcement Server (QES) for mitigation and quarantining of offending endpoints. This solution works in conjunction with the IP-assignment process, and can assign high-risk hosts to quarantine VLAN when they do not meet security standards and policies. The advantages and disadvantages of this approach were discussed in Part 1 of this paper. The most notable disadvantages of this approach are twofold:

- it is easily bypassed by using a static IP address
- it is not particularly effective at providing quarantine capabilities post-entry without the addition of yet another client on the endpoints themselves.

B. Switch integration

One of the most common quarantine capabilities uses the switch for mitigation, which provides two options. In the first, an ACL is used to block a particular device from the network and may drive it to a quarantine server. In the second, 802.1x is used to change the VLAN designation of a given endpoint to move it into a quarantine zone, which grants limited access and allows the device to connect to a quarantine server that resides in that zone.

Of these two, the first could be characterized as a brute force, "all or nothing" approach, which can have a negative impact on switch performance. The most common use of ACLs is simply to block a device's access, without providing any redirection to enable remediation. In this instance, both the user experience is impacted, as are costs associated with help desk activity, because the user does not know why his connectivity was terminated. Another drawback of this approach is that there is often a non-trivial lag between the detection of a threat or policy violation and the application of the ACL, losing critical seconds as the threat spreads.

C. 802.1x

802.1x typically doesn't impact performance as greatly as ACLs, but these solutions come with their own set of challenges. When coming up on an offending device, they typically either turn off a switch port that the offending device was connected to, or they put the offending device in a separate VLAN.

In the first instance, we see the same help desk and user experience problems as in the ACL scenario. Additionally, if the organization is using IP Telephony, it is very likely that the user's phone might also be off the same switch port. In this case, the user finds himself in the unenviable position of having an undiagnosed problem he can't report.

If, on the other hand, the offending device is moved to a separate VLAN, the IT manager should avoid using a broadcast VLAN, which is what most 802.1x solutions use. In the instance that the post-admission NAC technology discovers a threat outbreak, sending infected and at risk endpoints to the same VLAN results in more infections and drives up recovery costs.

But often, the 802.1x decision is easily made: a preponderance of installed networking equipment simply does not support 802.1x. Many organizations have hesitated to roll out 802.1x capability because of the complexity required in managing configuration on the switches and configurations on the endpoints themselves.

D. In-Line Blocking

For in-line post-admission NAC solutions, mitigation usually takes the form of blocking traffic at the in-line device, in combination with a Web redirect. In contrast to the ACL approach, this solution delivers improved response time and the granularity of mitigation. The solution can drop the traffic from the offending device only, rather than that of the whole network segment on which that device resides.

In-line devices' primary weakness lies in the fact that they don't see downstream traffic, and therefore cannot quarantine endpoints from their peers on the same broadcast LAN or other LANs where the traffic does not flow through the in-line device. This means that the higher up the in-line device is in the network, the more risk there is of broad internal infection in the instance of a threat outbreak because the threatening traffic may never make it to the security device.

E. ARP-Based Mitigation

This approach, which is used by Mirage Network's NAC solution, specifically addresses some of the drawbacks of other mitigation alternatives.

- First, it doesn't require any integration with the networking infrastructure, as with 802.1x mitigation, nor any other security server resource, as with DHCP mitigation.
- Second, it does not suffer from the problem of cross-infection that can occur when quarantine VLANs are employed. With ARP based mitigation, all quarantined devices are isolated from each other, and are only allowed to send traffic to designated quarantine servers or other approved devices.
- Third, it can look for and redirect all Web traffic to a specified server for user interaction, resulting reducing of help desk costs that can often come if the port is simply turned off.

1. ARP Management in a Nutshell

Essentially, ARP mitigation changes the ARP table on the offending endpoint, and directs all future traffic from that device through the NAC solution. In short, when Mirage NAC implements this approach, it becomes the gateway for the device in question and all communication to and from this device now goes through the Mirage NAC appliance. This approach is unique to Mirage Networks and is part of this company's core intellectual property.

II. Diagnosis and Remediation Tools

If the solution selected offers quarantining capability that includes redirection to a quarantine server, then there are usually a few things that that server will be used for: communication to the end user, form-based interaction with the end user, such as for registration, diagnosis of problems, which may include deeper vulnerability and OS scans or AV scans, and finally remediation of problems - cleaning of infections, patching of the OS, termination of disallowed services like peer-to-peer, to name a few.

The key aspect of the quarantine server is that the pre- and post-admission NAC solution can direct offenders to a device that can offer a variety of services, to the end users, to help them reduce the time and pain associated with getting their device in compliance with security standards and back on the network. Productive users are happy users, and happy users require less IT and help desk resources. This is a concrete ROI that is provided by NAC solutions and can greatly improve the efficiency of IT and help desk organizations in dealing with these issues.

IV. What's Next?

When we started this report, we set out to answer two critical questions IT security professionals find themselves asking with increasing frequency and urgency:

How do I control the access to my corporate networking resources?

-and-

How do I ensure that the resources that are allowed on my network aren't creating a security risk?

To answer those questions, we began with an examination of the factors that catalyzed the development of security technology, from the first antivirus technology in the 1980s to the rise of NAC as a network security hot button today. We then introduced the three core elements of a successful NAC implementation: pre-admission, post-admission, and quarantine and remediation. In Parts 2 – 4, we discussed those core elements in depth, defining alternative technological approaches and the pros and cons attached to each.

We've learned that no NAC solution – indeed, no security solution – is a 100% fix for all security problems, but that Mirage NAC can help IT pros answer those two questions with confidence.

Learn about the Mirage Networks NAC solution

Call, email or visit Mirage Networks at:
866.869.6767
info@miragenetworks.com
www.miragenetworks.com

About the Author

Scott Olson, CISSP

Scott is a 14-year veteran of computer security, and a recognized thought leader on the subject. He has had articles published in such publications as Computerworld and has spoken at security events such as RSA and Digital ID World. Currently vice president of product management and marketing at Mirage Networks, he is responsible for developing and leading product strategy and marketing initiatives. Previous experience includes executive leadership positions with Whole Security, a behavioral endpoint security solution provider acquired by Symantec in 2005; NetSolve, a provider of remote network management and security services, which Cisco acquired in 2004; and WheelGroup, which he co-founded to produce the first commercially viable Intrusion Detection System (IDS), and which was acquired by Cisco in 1998. Scott started his career as an early member of the Air Force Computer Emergency Response Team (CERT), where he tracked hackers and protected, tested, and monitored Air Force sites. Scott is a Certified Information Systems Security Professional (CISSP).