NAC DEPLOYMENT
A FIVE STEP METHODOLOGY

JOEL SNYDER
OPUS ONE

FEBRUARY 2007

OPUS one

Abstract

Deployment of Network Access Control (NAC) technology throughout the enterprise is a complex and expensive process. As with any IT project, the success or failure of a NAC deployment will depend, to a great extent, on the design and architecture development processes that take place well before the actual installation begins. This white paper offers a five-step methodology that will position any enterprise for achieving success with its network access control deployment.

Executive Summary

Adding Network Access Control (NAC) to an existing network is a dramatic and significant change to the physical network. When NAC is in place, the network is no longer a neutral substrate for moving packets around as quickly as possible. Instead, it becomes a security barrier; authenticating users, evaluating the security of end-point systems, and applying access controls focused on the user and their security status. A NAC-enabled network is no longer a utility, like power and water, but must be tailored to fit organizationally into networking, security, and desktop management teams to be effective.

This white paper discusses five critical questions that must be answered at the very early stages of any NAC project. These technology-independent questions form the basis of a deployment methodology. By addressing these questions before you've picked products or even chosen the IT team members who will be assigned to complete the project, it is very likely that you'll be able to address the most significant issues your team may encounter along the way to NAC success.

The five questions are:

1 | **What are your goals for bringing NAC into your network?**

2 | **How will you utilize user authentication within your NAC policy?**

3 | **How will you tie the end point security (also referred to as Posture Assessment) into your NAC policy?**

4 | **Where in your network will you enforce access controls, and how granular will your enforcement be?**

5 | **How will you ensure that your NAC deployment will be implemented systematically across your organization without causing unnecessary interruptions to your existing network?**

What is NAC?

The high-pitched buzz surrounding NAC complicates the matter of defining it simply. At its core, NAC combines user authentication, end-point security assessment, and access control. You're probably doing all of these things in one way or another. With NAC, these three elements are combined into a single solution.

NAC is user-focused, network-based access control.

The implications of that statement are fairly significant. Breaking it down will provide a clearer definition.

User Focused

The term "user-focused" differentiates NAC from many other forms of access control, such as that which you might find in a typical firewall. While a firewall provides access control, most are designed to be destination focused: it's not who you are, but what you want to get to that is most important in the decision to let you in or not. When "who you are" is considered in a typical firewall context, it's generally just a question of what your IP address is. NAC is different because NAC is focused on the user, and defines security policy (at least partially) based on the user's identity.

At a minimum, "user-focused" implies that users are authenticated and authorized. In other words, the policy expressed in NAC (the user's authorization) is based on who the user is as determined by some authentication mechanism. A full-fledged NAC solution might feature multiple authentication methods, including technologies such as 802.1X, captive portals, MAC-based authentication, and port-based authentication.

The second part of "user-focused" in a NAC scheme is the ability to include information about the user's device, along with the user's identity, in the overall access control policy. User-focused platform assessment requires a security posture evaluation of the user's access device. The most common approach to achieve this is to run some software on the user's device that reports the security status of the device, such as whether virus software is up to date and enabled. However, other approaches, including external scans, are also fairly common.

End-point posture assessment is just one piece of a larger part of user-focused access controls called "environmental" information, which also includes other data the NAC solution gathers from the environment, such as the access method (wireless, wired, VPN, for example), access location, type of device, or time of day.

Network Based

The second phrase in this definition, "network-based," means any NAC solution must sit in the network itself. Enforcement could be at the point of entry (at the switch or VPN device, for example) or it could be deeper (at a firewall or security device between the edge and the core), but NAC requires that the enforcement be within the network. It cannot take place on the client or at the end host (although other NAC components, such as end-point security assessment, will often run on the client).

Access Control

Finally, the term "access control" here means that you are restricting what hosts and services that end system can get to subject to the policy of the network manager, the authentication of the user, and the security posture of the end point. NAC can have many different levels of granularity and can even include combinations of technologies working together. The most common enforcement mechanisms are go/no-go access, VLAN-based access controls, simple packet filters, and full stateful firewalling. These controls could be aimed at controlling access to resources, or they could be used in a more primitive way, such as to simply enforce remediation of non-compliant end points.

Elements of a NAC Deployment Methodology

The most important parts of a NAC project are understanding, before any changes are made to the network, the answers to the five questions in the table below.

The remainder of this paper dives deep into each of these questions. While answering these questions taken individually won't represent a methodology for implementation, the collective answers will indeed pave the way to a successful deployment. In any case, if you can't answer these questions before you employ a NAC scheme, it's almost certain that deployment will fail.

Area	Key Question	Details to be Answered
Policy	What is your security policy?	What are you trying to accomplish? What type of users (such as guests or employees) and devices (such as mobile devices or company-owned laptops) will this NAC deployment focus on?
Authentication	What authentication method will you use?	How will user identity affect security policy and access control? How will you handle "failure" cases?
Environment	What end-point security (Posture Assessment) features do you want?	What types of devices will have their posture checked? What is the associated policy? How will you handle users and devices that cannot be checked, such as guests or printers? Will you be running continual posture checks, or just at login time?
Enforcement	What enforcement strategy will you use?	Where in the network will you enforce? Will you mix different types of enforcement, or use a single consistent strategy?
Integration	How is NAC going to integrate into your existing network?	How will physical integration be done? What steps can you take to ensure that integration goes smoothly and without unnecessary disruption? How will it integrate organizationally?

Question 1: What are your goals for NAC in your network?

Everything in the world of security is designed to reduce risk. Technologies such as encryption reduce the risk that private data will be made public. Firewalls reduce the risk that someone will connect to an application which they are not entitled to. Anti-virus software reduces the risk of malware infections.

With NAC, you must decide exactly what risk you are trying to reduce, and how much you want to reduce it.

When enterprise networks act as high-speed thoroughfares for information flowing through the enterprise, three critical resources may need to be protected: the network applications, resources, and the servers they run on; the users' machines on the network, and the integrity of the network itself.

An early step in setting NAC policy is identifying which of these three resources is your protection priority. Most commonly, enterprises will start a NAC project with a focus on critical applications. However, other enterprises will chose a NAC path for other reasons entirely, such as regulatory compliance.

The following list will help determine the risks most relevant to your NAC deployment:

A | Are you trying to keep malware off your network?
In this case, you're probably not as worried about who is connected as much as you are worried about what hardware and software are connecting to the network. The risk is that of a user who knowingly or inadvertently connects with an infected system that can pass the infection elsewhere in the network. If this is a key risk, your focus should be on end-point security assessment, both using downloadable tools and even external auditing and scanning techniques, combined with intrusion prevention and detection systems.

B | Are you trying to help honest people stay honest, by checking their adherence to your security policy?
In this case, you should be most concerned about compliance, although obviously one of the goals of the security policy is normally keeping malware off the network. With honest users, you may not need to be as concerned about someone intentionally trying to subvert the NAC infrastructure. If failure to comply with policy is a key risk, the focus should be on both end-point security assessment and, more importantly, remediation.

C | Are you trying to keep attackers off your network?
In this case, you're typically facing a situation where wired and wireless networks—some of which may even be located in very public places—are accessible to untrusted individuals. The risk in this case would be a dedicated attacker less interested in passing on malware but more interested in stealing your data or using your applications on your servers. If this is a key risk, your focus should be on authentication mechanisms and access control policy, coupled with intrusion prevention technologies, as part of your NAC scheme.

D | Are you trying to take greater control over network access in general?
In this case, you may have a largely uncontrolled network from a security point of view that needs additional boundaries and points of control to keep different groups where they belong on the network. While traditional techniques, such as putting developers in one building and administrative types in another, can provide physical segmentation, these techniques can be augmented with NAC. Reducing risk in this instance means pushing the control point closer to the user, and regulating that access based on who the user is, rather than where they are physically located. If greater control is key, your focus should be on cost-effective and non-intrusive authentication mechanisms and access control (such as 802.1X authentication and VLAN-based access controls), perhaps trading off enforcement granularity so as to minimize disruption.

E | Are you trying to track network use and in the process prevent misuse of your network?
In this case, the actual access control may not be as critical as passing traffic through devices that can map an authenticated user to their traffic—and then react when something is not quite right. The risk here may be a compliance and regulatory one, rather than a traditional security concern. If this is a key risk, your focus should be on authentication as a critical step to gaining network access, followed by sufficient instrumentation to enable you to map the session back to the authenticated user—a distant reality in many NAC products with poor reporting and session tracking capabilities.

F | Are you trying to balance guest and employee access simultaneously?
If you're primarily worried about security threats from insiders, that will take you down a different path than if you're worried about threats from guests who may be temporarily using your network. The risks that come from each type of user are similar, but your strategy in dealing with them will be very different. If guest users are a key risk, your focus should be on access control that allows you to identify guest users and enforce a strict security policy that keeps them off most or all of your internal network with tight firewall controls. On the other hand, if employee users are a key risk because they might be

using applications or servers inappropriately (or even accidentally), your focus should be on authentication, with simpler access controls, such as VLAN-based restrictions.

The facile answer might be "I want to do all those things." If that is the case, you still have to prioritize what types of risk you care about most.

Resolving Risks by Deploying NAC—Define Your Priorities

1 | Are you trying to keep malware off your network?
2 | Are you trying to help honest people stay honest?
3 | Are you trying to keep attackers off your network?
4 | Are you looking for greater control over network access?
5 | Are you trying to track network use and prevent misuse?
6 | Are your threats coming from guests or employees?

Once you've identified what risks you're trying to reduce, and their relative influence on your network, you can start assessing the components of NAC and prioritizing which components get deployed into your network.

Any NAC solution is going to bring three components to your security toolbox:

1 | User authentication
2 | End-point security assessment (and remediation)
3 | Access control enforcement

By identifying the pertinent network risks, you'll have a good idea of which of these three components will play the biggest role in your NAC deployment and how each will have to be configured to meet your risk reduction requirements. For example, you may have determined that access control enforcement is key to your NAC business case, and that highly granular enforcement is necessary.

After you've identified your risks and prioritized which NAC component will help reduce those risks, the last step in defining your NAC goals is to decide where in your network NAC will apply. Many security analysts blindly assume that NAC is everywhere, all the time. But in reality, NAC may well be better suited as a means to solving a particular problem, such as guest access in public areas or insecure branch offices. It's certainly true that you will best be able to leverage an investment in NAC if you apply the technology broadly, but there has to be a good business reason to start using NAC in each part of the network

Use the table entitled "Where will you deploy NAC?" as a starting point to identify areas of your network that can be protected by NAC. In each of these areas, decide based on the risks you identified whether NAC is appropriate or not.

With your explicit list of risks, a prioritized set of NAC components, and knowledge of where NAC will go in your network, you're ready to refine your strategy by looking into each of the components in greater detail.

Where will you deploy NAC?

A | To allow VPN remote access
B | To support local wireless users
C | To permit guest wired and wireless users (e.g., public areas, conference rooms)
D | To check local end-user desktops and laptops
E | To control remote and branch office access
F | To protect infrastructure by enforcing access controls on servers and other embedded devices in the data center.

Being that NAC is user-focused security, you have to know who the users are. Any NAC deployment must employ some set of authentication methods. Three dominant methods have emerged : 802.1X-based authentication, web-based authentication, and proprietary-client authentication. While some NAC vendors have additionally proposed "passive" types of authentication (such as watching a Windows login sequence fly by), these approaches vary wildly, offer a lower level of security, and have not yet been widely tested.

Since authentication is one of the two critical factors that will determine access control policy (the other send-point security assessment), you have to assess what authentication method you will use.

Three Common Methods of Authentication

The most secure authentication method for NAC is based on 802.1X, the IEEE standard for authentication over local area networks. In an 802.1X environment, the user's device is not connected to the network – wired or wireless – until the authentication is successful. Without authentication, the user doesn't get an IP address, can't sniff traffic, and certainly can't attack anyone or anything else on the network.

A positive side-effect of 802.1X authentication is the creation of an encrypted tunnel between the end user's computer and the authentication server inside of the network. Several of the popular NAC frameworks, including the Trusted Computing Group's Trusted Network Connect protocols and Cisco's NAC framework, leverage this encrypted tunnel to bind authentication and end-point security assessment together.

Question 2: How will you handle authentication?

Authentication based on 802.1X is well-supported by software built-in to both Windows and Macintosh desktop and server operating systems, and by all modern wired managed LAN switches and enterprise-oriented wireless access points. Indeed, the 802.1X authentication tunnel is also used in wireless environments to provide encryption keys as used in the popular WPA and WPA2 (802.11i) wireless security standards.

The core authentication system inside of 802.1X is the Extensible Authentication Protocol (EAP), which can also be extended to other environments, such as VPN authentication in IKEv2.

A different approach to NAC authentication can be found in web-based authentication methods. With web-based authentication, the familiar "captive portal" model is used to gather authentication information from the user. (A captive portal is the mechanism used in environments such as public hot spots, where the user is redirected to a web page after they launch their browser. This page then prompts for authentication or payment information.) While much less secure than 802.1X-based authentication, because by the time you get to the captive portal you already have an IP address, have been able to sniff traffic, and have gone at least a few places on the network, web-based authentication benefits from the wide availability of web browsers and the easy familiarity of end users with the idea of entering a username and password (or other authentication method) on a web page.

In web-based authentication, a user connects to the network, is given an IP address, and has some level of network visibility and access. If the user opens a web browser and attempts to view a web page, the NAC solution, through a variety of techniques, intercepts the web request and redirects the user to a page asking for authentication information. As with 802.1X authentication, this captive portal process doesn't imply any particular authentication method: username and password, token card, or any type of common authentication works well with web-based captive portals.

A benefit of web-based authentication is that it dramatically simplifies the process of downloading and installing software (such as end-point security posture assessment tools) on the end-users' systems. Because the user already has a browser opened up to a web page, stuffing additional software down that path is easy to do and, again, benefits from the easy familiarity of end-users with their web browsers.

Finally, proprietary client authentication offers a less-desirable alternative to either 802.1X or web-based authentication. While proprietary methods are going to vary from vendor to vendor, the common elements include some sort of installed client software on the end-users system and a greater topology flexibility than 802.1X provides. Any proprietary client method also requires a consideration of how this will be installed and licensed on devices that don't have it, as well as on unmanaged guest devices.

For example, 802.1X generally requires that the user be authenticated at the point of entry to the network, at the edge switch. With proprietary client authentication, a device closer to the core such as a firewall or router can both control the user's access and act as the point of authentication.

Since many NAC deployments include end-point security assessment as part of the access control decision, a proprietary client can greatly ease the task of actually performing the assessment. Since the client is resident and installed on the end-user's system, it can serve the dual tasks of authentication and end-point security assessment without requiring the user to download additional software or allow it to be pushed through a browser connection.

PROS AND CONS OF NAC AUTHENTICATION METHODS

	802.1X	Web-based	Proprietary Client
PROS	Highest security; standards-based; multiprotocol; most transparent; scales; built-into modern operating systems.	Very familiar model to end-users; broadest platform support; handles guest users best.	Tight integration between client and security policy; broad range of topology support.
CONS	802.1X supplicants have a "bad reputation" (although this is not supported in our testing); weak guest support; poor support for non-mainstream platforms such as Linux, Palm, Symbian, and embedded devices.	Onerous and slow for all users; only supports IP; requires web browser; security model weaker.	Platform support not broad (usually Windows-only); requires vendor lock-in; weak guest support.

Balancing Authentication Methods

Each of the authentication methods discussed above has pros and cons. The table above summarizes some of the most significant ones.

The best help you'll have in selecting an authentication method is your own policy. Some network managers won't care about authentication, because it's not part of their risk-reduction strategy. If you do have authentication as a priority, which most NAC deployments will, your policy and goals will help you select the most appropriate method.

For example, if you're focused on enterprise users and are implementing NAC to help keep enterprise desktops in compliance, both on the LAN and through a VPN, you'll want to use technologies such as 802.1X. On the other hand, if you're looking to NAC to fix the problem of giving guests access to your network or occasional staff use (such as in conference rooms), then web-based authentication is a better choice.

One smart strategy to consider is instituting a fallback alternative authentication method when users don't have the software (802.1X, browser, or proprietary client) needed to fit into your NAC strategy. For example, a maximum security strategy uses 802.1X for authentication. Unfortunately, that doesn't handle guests and visitors very well who might not be using 802.1X or have credentials on your network. With a well-designed NAC solution, 802.1X could "fallback" to a different authentication method such as web-based authentication if the network infrastructure detects that 802.1X is not supported on the client system—ideal for guest users.

Finally, your authentication strategy must accommodate the growing set of devices that won't speak 802.1X, don't have a web browser, and can't run a proprietary client. These are often mobile devices and represent a class that will only grow with the continuing spread of Ethernet and wireless connectivity to anything with a battery. Your authentication method choices for such devices are severely restricted, and may be limited to weak authentication using MAC addresses. In that case, the basic authentication of the device should be accompanied by a very strict access control policy and stateful firewalling to limit access to the small set of services and servers that are absolutely required.

Question 3: What end-point security policy will you enforce?

End-point security is a dichotomy in that it's both the killer application that may drive NAC adoption as well as the Achilles heel of NAC because it works best when you need it least.

End-point security assessment is designed to answer the question: does this system comply with the security policy of the enterprise? The end game of combining a security policy with compliance assessment is risk reduction. It is the goal of risk reduction, rather than the simply policy compliance, that should drive your NAC deployment strategy.

Reducing Risk with End-Point Security Assessment
You will find that when designing a NAC deployment you are constantly examining the edge cases where a system or user does not comply with policy, or (more commonly) where you cannot determine the level of compliance. In these cases, your design should focus on balancing the risk of letting a non-compliant system onto the network against the cost of denying access to a legitimate user. For example, you will want to be stricter about policy compliance when especially sensitive and valuable resources are involved than you would be with a user who only has rights to the Internet or an internal webmail server.

Defining end-point security assessment policy, and deciding how to handle the inevitable failure cases, is one of the most difficult parts of planning a NAC deployment. When your NAC deployment is solely focused on employees and managed desktops or laptops, the task is easier. But as NAC spills out into the realms of guest users, casual users, or users with their own computing devices, building a credible policy that actually reduces risk becomes a delicate balance between security and usability.

One danger with end-point security (EPS) assessment is that every tool created by every vendor has some failure rate in the form of false positives: the tool says that a device doesn't comply to policy, when it really does. End-point security assessment tools will also have false negatives, reporting that a system does comply to policy when it really doesn't. These are less common, but are a cause for caution, especially if deliberate deception is suspected. In these cases, your NAC strategy should have checks and balances, such as intrusion detection systems, that can help to crosscheck end-point security assessment. EPS assessment tools may generate false positives when they're incompatible with end user platforms, or just because they can't get the information they need.

End-point security assessment has other facets that can further complicate deployments. Some NAC solutions include tools such as external auditing systems or data feeds from IDS and IPS devices. For example, an IPS could notify the NAC infrastructure that a new client is generating an unusually large number of alerts, in which case the NAC infrastructure could tighten access controls or even block the user entirely. Similarly, the NAC infrastructure could notify an IPS that a client about to be admitted to the network is less trusted (perhaps because it could not run EPS assessment) and the IPS should apply a higher level of protection to traffic from this client. These can help to reduce the uncertainty of the answers that EPS provides, although at the cost of increasing complexity.

The potential for end user frustration, aggravation, and annoyance at in the name of end-point security must be factored into the policy and deployment. As a new technology, NAC is subject to the same requirements of every other new technology: NAC will be adopted to the extent that the pain it causes the end users and IT staff is less than the pain and risks it reduces. A clear danger of any EPS policy is the potential to swing the balance in such a way that NAC creates more problems than it solves.

While there are no clear formulas for calculating this balance, you can predict that a very homogeneous network with tightly managed desktops and laptops will have a high success rate with EPS assessment, while one where systems are heterogeneous and not centrally managed will have a higher failure rate.

An iterative process of testing and refining end-point security assessment will help here. You will want to be flexible enough in your planning to ensure that the burden of EPS is commensurate with the level of risk reduction.

Continuous Assessment and Remediation

Remediation is an important part of the end-point security assessment component of your NAC design. Obviously, if a user's system is non-compliant, there are significant benefits to helping solve the compliance problem, rather than simply shutting the door on the user and dumping them in the street. Remediation strategies can range from simple one-size-fits-all methods to more fine-grained approaches that apply different assessment and remediation strategies depending on other variables.

A common example of NAC-facilitated remediation is anti-virus policy compliance. If the policy says that a user's system must be current in its anti-virus signatures, then giving the user access to the signature files so that they can download them and become compliant is a good idea.

A slightly different example might be anti-virus scanning. If policy dictates that anti-virus scans must be run every 24 hours and the user's system been turned off for a month, helping the user understand how to launch a scan or re-enable the anti-virus software would help remediate the problem and put the system into compliance.

Finally, you may want to consider auto-remediation as part of a NAC deployment. If the policy requires that a personal firewall be turned on but it's not, your end-point security assessment software could simply attempt to turn the firewall back on, bringing the system into compliance.

All types of remediation—granting access to resources, providing assistance and information, and auto-remediation—are aspects of end-point security assessment that you should consider in your NAC deployment. You may also have reasons to reject self-remediation and auto-remediation. For example, your policy might call for a more detailed examination of a system that has fallen out of compliance.

While providing remediation resources adds complexity to the NAC deployment, the value of remediation in facilitating network access and avoiding user frustration and service denial may make the effort and cost worthwhile—not to mention the money that will be saved on the help desk

In some threat scenarios, deliberate deception is a real risk. The most common case would be a user who turns on anti-virus software to pass the EPS assessment, but then turns it off after access is granted. In these cases, policy might call for continuous enforcement, checking and re-checking the status of the system to be sure that compliance is equally continuous. Be careful here, as the burden of continuous EPS assessment should be balanced against the risks involved. Some NAC tools have absurd requirements in order to provide continuous enforcement, such as asking the user to keep a particular web browser window open at all times. If continuous enforcement is needed in your environment, be sure to evaluate the user experience in detail.

Other deception risks are more sinister. For example, a system requesting access may be so heavily compromised that the end-point security assessment tools might not be trustworthy, and the tools might "lie" to the NAC servers about the state of compliance, claiming that a system is clean and compliant when it isn't at all. This is a difficult problem to solve. The Trusted Computing Group (of which Trusted Network Connect, TNC, is a subgroup) is specifically trying to address this problem in a generic way with its Trusted Platform Module (TPM), which supports the notion of a truly trustworthy software module resident on the client. If you are concerned with this type of deliberate deception as a threat, then your NAC deployment must address this concern.

Question 4: How will you enforce access controls in NAC?

Access control enforcement is the final result of all the authentication and end-point security assessment your network has just completed and your users have just endured. Once your network knows who the user is, and has discovered the status of their end point, then it can enforce access controls.

Enforcement in NAC is closely tied to specific network topology and overall network capabilities. While some enterprises might have the luxury of a full-scale replacement or augmentation of their existing network, most will need to compromise between what an ideal security policy might call for and what existing hardware allows. The question for many NAC deployment planners is "What can I do with the hardware I already have?" Which is closely followed by "How can I achieve my goals with a minimum of additional investment and disruption?"

NAC itself doesn't pre-suppose a particular enforcement technology or even strategy, for that matter. Four common options, ranging from simple go/no-go, through VLAN assignment, packet filters, and up to full stateful firewalling, will cover the majority of deployments. As with many aspects of a NAC deployment, you don't have to select a "one size fits all" approach. For example, using VLANs, guest users might be shunted off to a lightly firewalled network that gives them Internet access only, while internal users might be subjected to a combination of technologies, including VLANs, internal firewalls, and intrusion prevention systems, all acting in concert to varied enforce access controls.

NAC enforcement comes down to two main questions: "Which types of access controls will you use?" and, "Where in the network will you enforce control?"

These questions don't require single word answers. In fact, there are excellent reasons to mix and match enforcement methods during the life of your NAC deployment. For example, you may want to start with a simple access control method to gain experience with NAC, and then add in a more powerful one as you become more comfortable with the technology. You might also want to combine methods to achieve a more traditional defense-in-depth strategy. For example, you can use VLAN-based access controls to coarsely separate users, while adding NAC-controlled stateful firewalling deeper in the network to protect more critical assets.

Access Control Options

Most NAC vendors are offering one (or more) of four different options. In order of increasing security, they are: a go/no-go to network access; VLAN assignment; basic packet filters; and, a full stateful firewall.

Obviously, full stateful firewalling of every single user is what a security manager who could have everything would ask for. However, not only is pushing firewall technology all the way to the end-user's port very expensive, it also puts you on the bleeding edge of network technologies because there are very few firewalls currently available with very high port density and high performance. Achieving this security nirvana of full control would require massive changes in infrastructure.

However, there are intermediate steps that give many of the benefits of firewalling without the same costs and bleeding-edge danger. Typically, you accomplish this by putting coordinated, very high-speed firewalls deeper in the network, and controlling the firewall rule sets with a NAC policy server. A number of firewall vendors focusing on NAC, both established and start-up, are offering products in this space.

Basic packet filters (generally called access control lists (ACLs), don't offer stateful firewalling, but are closer to the existing capabilities found in many installed switches. Using basic packet filters may let you push enforcement to the point of access of the user, such as a switch port, which does have advantages. As a compromise approach, there's some distance between stateful firewalls and basic packet filters.

While firewalls specialize in large and complex rule sets, switches often have significant limitations on how packet filters can be loaded, how many can be loaded, and how long they can be. This means that some generality available in full firewalls (such as the ability to combine different group memberships to form a "super-rule") might not be available in basic packet

filters. Nevertheless, this type of technology predates the term NAC by several years and is broadly available, especially in single-vendor switch deployments.

VLAN-based enforcement is the most commonly mentioned NAC enforcement strategy, even if the granularity is quite coarse. With VLAN-based enforcement, a user is placed on a particular VLAN based on their identity and end-point security status. It's assumed that static firewalls placed between VLANs are actually controlling access.

With VLAN-based enforcement, users must be divided up into very coarse groups. A typical deployment might have between five and fifteen different VLANs, including one for segregating guests, one used to quarantine users for remediation purposes, some designated for dedicated devices (such as printers and VoIP phones), and a small number used to differentiate internal users. Creating more than a handful of groupings using VLANs can quickly become unmanageable, especially in a large campus environment.

VLAN-based access controls have two other issues that require attention. First, using VLANs as security barriers means that the management of the network switching infrastructure is now as important as the management of your firewalls. In fact, it becomes the same thing. If any single switch in the network is compromised, then so is the entire security of the network. Since network management teams have not traditionally treated switches as security devices, using VLAN-based access controls alone will require greater attention to issues such as switch firmware updates, along with a sharing of configuration between network and security teams.

Second, you still have to depend on additional firewalls to both route packets and mediate the access between VLANs. These external firewalls have to be specified to have sufficient performance to handle LAN-to-LAN routing, and sufficient availability and scalability to be installed in the core of a network. And, because the firewalls are now part of the bigger picture of NAC enforcement, the firewall policy and rule base have to both accommodate traditional firewall functions as well as NAC functions, again drawing together the network and security teams. VLAN-based enforcement fits into a network where the set of different security policies is very small, but is not as simple as a first glance might indicate.

The weakest of the access control strategies is "go/no-go" enforcement. With "go/no-go" access controls, users are either allowed on the network or they are blocked, based on successful authentication and end-point security assessment. It is a model that has worked well in many areas, such as wireless LAN access where the switch only needs to verify whether the device knows the WEP key or WPA secret and can get on the network, or doesn't, and it's locked out.

The "go/no-go" access control strategy has its greatest value as a stepping stone to more comprehensive NAC solutions. For example, you might want to deploy end-point security assessment first, and use that as a gating decision for access control. That would give you experience in EPS that can then be leveraged into tighter access controls when authentication is added. Of course, "go/no-go" based on EPS without remediation carries its own dangers—so when rolling out a "go/no-go" solution, keeping careful watch of logs and help desk calls may be the only way to find out when things are going wrong. One thing is clear: jumping from no NAC to full NAC overnight is an unwise choice, and "go/no-go" as an intermediate strategy can reduce the risk of making an expensive deployment error.

Question 5: How will NAC integrate into your existing infrastructure?

No one gets to start from scratch when it comes to NAC, and that means that any NAC deployment must merge into an existing infrastructure. In this case, infrastructure doesn't just mean the hardware you've already got, but also the software, policies, procedures, and even organizational infrastructure.

Some NAC deployments will be small in scope and easy to add onto an existing network. For example, if you're simply adding NAC features for guest access to an outward facing wireless network and the main conference room, this should not be a dramatic change to the network. However, in a full-scale enterprise deployment, NAC represents a significant and dramatic change to the organization. Adding an entire layer of security to an existing network is as monumental a change as installing the network in the first place. If you do not consider how an enterprise-scoped NAC deployment will integrate into the hardware, software, and policies of the organization, you put the deployment at significant risk of failure.

Integrating into the Organization

NAC is unique in that it will require tight and intense cooperation between three different enterprise IT teams: the team responsible for the network, the team handling security, and the team managing the desktops which is also knowledgeable about Windows security. In some organizations, these teams are well-integrated, but in others they may not even inhabit the same building or have the same reporting lines. Because NAC touches each of these disciplines, getting members from each of the teams together early and often is important.

The enterprise itself may gain value from a NAC deployment, but this doesn't mean that the enlightened self-interest of each part of the organization shouldn't also be served by NAC. It is fortuitous, then, that each part of the organization stands to benefit from the addition of NAC.

The Windows-savvy desktop team will gain additional enforcement teeth for their security policy, as users who are non-compliant will be blocked from network access. NAC-specific reporting and management tools will also help this team to be proactive in addressing user security problems and helping increase their service level.

The security team will gain a stronger set of tools for managing the flow of information around the organization. This can help in regulatory and industry compliance initiatives. Likewise, stronger access controls reduce the potential for security breaches. This team can spend less time groveling through forensics logs to understand an incident, and more time on refining the security controls that prevent incidents.

The network team will gain an increased ability to manage performance and increase availability of the network by having a more predictable and segmented population and fewer incidents. Since NAC will generally require more attention to be paid to network infrastructure all the way down to the desktop, networking teams can also extend their knowledge and configuration control of the network to its very ends.

Integrating into Physical Infrastructure

Infrastructure integration also, of course, means physical infrastructure. When deploying NAC, most enterprises will try and minimize disruption and expense by re-using as much of their existing hardware as possible. While that's an admirable goal, it's also important to see where existing hardware restricts capabilities and where simple replacements can optimize results. Therefore, a detailed inventory and network survey of any area to be touched by NAC is a requirement before any NAC deployment can begin.

Equipment re-use is likely to be a prime strategy as every organization seeks to preserve value and eliminate unnecessary capital expenditures. This suggests that there will be a tendency to re-use existing elements for security purposes, whether or not security was an original design goal of the element. A common example is the use of SNMP—an unreliable, unacknowledged protocol—for security management. SNMP's original design was for the collection of statistics, where lost data was not significant. Re-purposing SNMP controls from data collection to security alerting and management puts a square peg into a round hole. A better strategy is to ensure that the security management protocols are reliable and authenticated—encrypted using industry standards such as SSL/TLS for the underlying secured transport –even if that means changing devices or limiting what equipment can be used.

In NAC projects, the amount of time spent understanding the real physical and logical topology of the network, along with the different types of devices and their location, commonly far exceeds the actual time it takes to install NAC-specific hardware and software—by a factor of 10 or more. While you don't have to understand every last aspect of your topology and network to deploy NAC, there are many parts of NAC which simply won't work and can't be configured until you have a clear handle on how security and networking really works inside the organization.

Next Steps: Deployment and Best Practices

Once you've made a good start by answering the five questions outlined in this white paper, you'll probably be itching to start installing NAC equipment and software. The deployment of a NAC solution should spread from a test environment into a larger rollout in measured steps and with the same deliberation as any other change.

You may find that anchoring your NAC strategy into specific use cases—risks you are trying to reduce, or specific problems you are trying to solve—will help direct your deployment. For example, you may want to specifically solve the problem of network access in conference rooms. Or you may be worried about the next virus outbreak, in which case paying attention to end-point security assessment first may be the appropriate start.

When you have a technology already deployed, such as 802.1X-capable switches or an end-point security-aware SSL VPN, that can also be a good base for a NAC deployment. Since NAC is as much an amalgamation of technologies as it is a new idea, you may already have NAC running in your organization in one form or another already. For example, if you are using 802.1X for wireless security with WPA2/802.11i, you can add access control or end-point security assessment to build on technology you already know and have installed.

If your NAC solution offers it, you should also consider deployment in "auditing" mode (rather than full blocking mode) to see what would have happened before actually interfering with existing user access.

In the world of IT, "best practices" are really the collective wisdom amassed after thousands of projects—and no one can claim to have touched that many enterprise NAC projects.

However, here are some starting points from other IT disciplines that should overlap with NAC deployments:

A | Break down your deployment into tasks and subtasks so that you don't attempt the futile process of putting NAC across your entire network overnight.
When you do enable NAC, consider implementing in stages. For example, you might enable authentication only without turning on end-point security assessment to work out the problems one subsystem at a time.

B | Maximize your investment by extending NAC as far as you can.
Although you can push NAC out to solve a single "point problem," once you've begun a solid NAC deployment, the best way to drive down the cost of management and acquisition is to use that same technology in as many places as you can in your network.

C | Pay attention to all devices on your network
Obviously, the most important goal of NAC is to support the network access needs of the organization's primary users. You should spend most of your time and effort on the core users and your core reasons for installing NAC in the first place. However, if you ignore the less common cases, you run the risk of coming up with a solution that simply won't work over the long haul. Many networks have so many "unusual" devices, that the unusual ones far outnumber the "typical" case. Embedded devices of all kinds, as well as PDAs, WiFi devices, VoIP phones, and printers are all examples of places where business-critical networking happens without Windows or web browsers. It is certainly fine to simply acknowledge how your solution will work around these devices, but it is dangerous to just ignore them.

D | Pay attention to all users on your network
Large networks generally have a diverse user population using a wide variety of computing environments. For example, many networks have a significant population of staff-owned laptops and even desktops. Even when personal systems are not used in the building, remote access via VPNs from home and when traveling represents a network access scenario that should be covered by a NAC solution—but may need to be considered and dealt with separately.

E | Be mindful of growth
Even if you are starting with a limited rollout and constrained scope, your NAC solution may eventually have to support the entire network of an organization, including remote offices, data centers, and campus wired and wireless networks. While you don't want to be paralyzed by looking too far into the future to pick a solution that will work for 10 users today and 10,000 users tomorrow, you do want to keep the inevitable scope creep of any project in mind.

F | Build for reliability and scalability
As with any network-critical technology, NAC requires careful attention to both scalability considerations as well as reliability and high availability. Many NAC solutions available today are first-generation solutions built by companies inexperienced in enterprise networks and security. While new ideas and fresh blood are great ways to spur technology innovation, you want to constantly be aware of the potential for a NAC solution to hit a scalability wall, or to itself become a single point of failure in your network.