Whitepaper

Selecting An Approach For NAC Enforcement: Five Key Issues

Joel Snyder
Opus One

Network Access Control (NAC) is based on a simple idea: what you can do on a network is a function of who you are and the state of your end-point security. NAC is not a single product that you buy, but a set of technologies that are brought together—sometimes by a single vendor and sometimes in a multi-vendor architecture—to implement this underlying idea.

To build a NAC solution, you have to bring together three specific security components under a common management umbrella. The two starting components for NAC are authentication (identifying who the user is) and environmental information (identifying, among other things, the state of end-point security on the user’s device). The third, and most critical component, is enforcement: making sure the user does, in fact, only go where they have permission to based on NAC policy.

Throughout the spectrum of NAC products, vendors and system integrators have identified four common approaches to enforcement. This white paper discusses these methods and points out key issues related to selecting an approach for the large enterprise network.

September 19, 2007

CONTENTS

EXECUTIVE SUMMARY... 1

FOUR APPROACHES TO ENFORCEMENT 2

Edge Enforcement 2

In-line Enforcement.. 3

Hybrid Enforcement.. 3

Protocol-based Enforcement 4

SECURITY.. 5

FLEXIBILITY.. 8

PRAGMATIC DEPLOYMENT AND RISK OF FAILURE... 9

DISTRIBUTED ENFORCEMENT, PERFORMANCE, AND SCALABILITY.................................. 12

COST 13

CONCLUSIONS.. 15

About Opus One® 15

EXECUTIVE SUMMARY

"Enforcement" in the world of NAC is the Control part of Network Access Control. In the world of NAC products, there are four enterprise-class approaches to enforcement: edge enforcement, hybrid enforcement, in-line, and protocol enforcement.

Edge enforcement uses a device at the edge of the network, typically a switch, to enforce access controls. In-line differs from edge enforcement in that the device enforcing access is not at the edge, but deeper inside of the network. Hybrid Enforcement, common in early NAC products, combines both in-line and edge enforcement techniques. Protocol enforcement uses network-based protocols to limit access by altering layer 3 services.

Selecting an approach for NAC enforcement is an important early decision in any NAC deployment. There are five key criteria that can be used to help differentiate enforcement approaches and to identify the one that is most appropriate to any network, including security, flexibility, risk, scalability, and cost.

- Security: NAC is primarily a security service, and thus the most important criteria for selecting a NAC enforcement method is the security of the enforcement. Higher security comes in approaches that force authentication, restrict pre-authentication network access, and tightly bind the authentication dialog to a device. Edge-based enforcement offers the highest security based on these criteria.

- Flexibility: Every security architect desires the maximum flexibility in the products they pick to deploy. Getting locked into "one way of doing things" is dangerous, and limits the ability to respond either slowly or quickly to new threats or new business requirements. NAC enforcement approaches that have greater flexibility are more desirable, especially with regards to evolving security technology. Because edge and hybrid enforcement allow for go/no-go enforcement, VLAN-based enforcement, stateless access control lists, and stateful packet filters, they offers the greatest flexibility.

- Risk: Taking a pragmatic, step-by-step approach to deploying NAC really increases the chance of success. It is better to deploy a NAC solution in small pieces, merging it cleanly with the network, increasing the level of security and control as you grow comfortable with the reliability and stability of the solution. Edge and, to some extent, hybrid enforcing access controls have the lowest risk and the greatest chance of success because they allow easy port-by-port rollout of a NAC deployment.

- Scalability: A potential danger zone for NAC is scalability. If NAC stops working, so does everyone in the company. This means that the NAC solution should scale easily as the network itself scales. Edge enforcement, and to some extent hybrid enforcement, benefits greatly from the distributed nature of the enforcement mechanism. Enterprise-class switches are engineered for a particular load based on a number of ports and traffic levels, and don't represent a bottleneck when operating within the engineered limits. In-line devices, especially those based on general purpose computers, can easily become a bottleneck.

- Cost: NAC should leverage of value in equipment that is already deployed. For example, NAC should use the features that are already present in installed equipment, such as 802.1X authentication and VLAN capabilities. The other is to build on the strength of network infrastructure vendors, all actively looking to improve the security of their products. Pushing

enforcement to the edge is both cost effective (since no new enforcement hardware is needed) and leverages the security features that are being pushed into switches.

A wise security architect looks to provide high security, good flexibility, low risk, predictable scalability, and reasonable cost in the solutions they design. While every deployment scenario is different and enterprise requirements can vary tremendously, edge-based enforcement an excellent choice for most enterprises looking to add NAC into their existing networks.

FOUR APPROACHES TO ENFORCEMENT

"Enforcement" in the world of NAC is the Control part of Network Access Control. When the NAC server derives the access control policy for which parts of the network a user should be allowed to use, some part of the network has to enforce this decision. In the world of NAC products, there are four main approaches to enforcement. Most NAC products depend on a single strategy to enforce NAC decisions, although some NAC products can use a combination of approaches.

The NAC industry has not necessarily agreed on the terminology used for these approaches. The terms used in this paper are the most descriptive we could create to describe what is actually happening within each approach.

The four approaches to enforcement are Edge Enforcement, In-Line Enforcement, Hybrid Enforcement, and Protocol-based Enforcement.

Edge Enforcement

Edge Enforcement, the easiest approach to understand, means that the device at the very edge of the network (typically a wired or wireless switch, although this could also be a remote access VPN concentrator) is responsible for enforcing access controls. A common example of this would be a typical LAN switch with 802.1X (the IEEE standard for layer 2 authentication), although strategies such as MAC address-based authentication are quite common in early NAC deployments. With edge enforcement, users authenticate before they can send a single LAN packet. The enforcement of access controls is done at the exact point of attachment to the network. All modern switches wired and wireless switches are capable of edge enforcement using a range of controls, such as "go/no-go" or VLAN assignment.

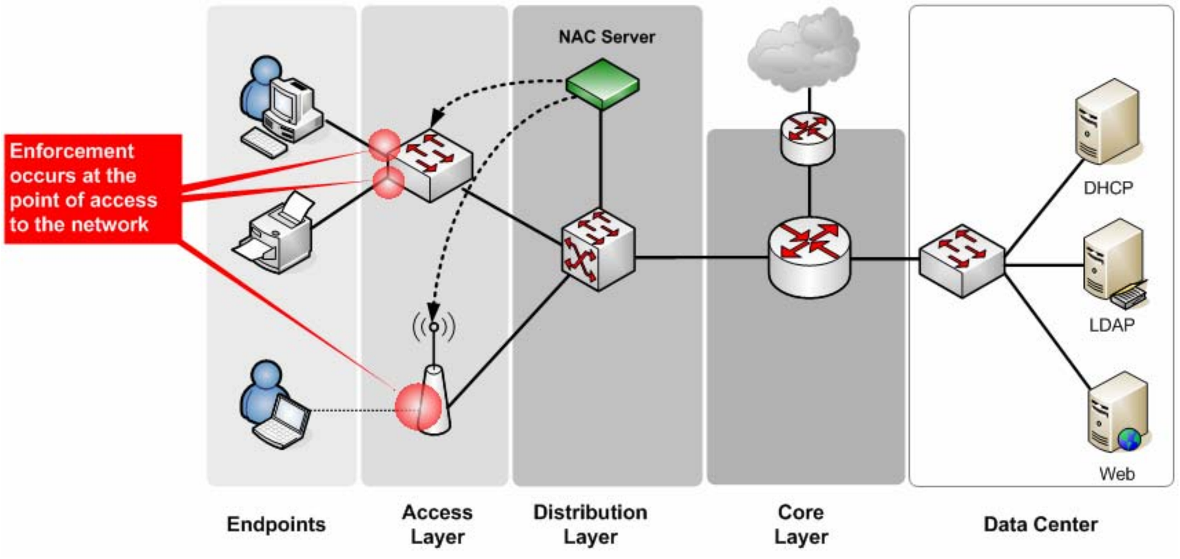

Figure 1- Edge Enforcement

In-line Enforcement

In-line Enforcement, also easy to understand, means that some device which sits between the user and a protected resource is responsible for enforcing access controls. In-line differs from edge enforcement in that the device enforcing access is not at the edge, but deeper inside of the network. The most common example of in-line enforcement occurs at a typical wireless hotspot, where some device is placed in-line between the wireless users and the rest of the network. In this example, users are forced to authenticate using the in-line device's web portal before they can go any further into the network with the in-line device then providing whatever access controls are appropriate, including both coarse access control rules and fine-grained stateful firewalling. Although in-line devices are often placed very close to the user, they can also be pushed much deeper into the network, sitting in front of the resources being protected.

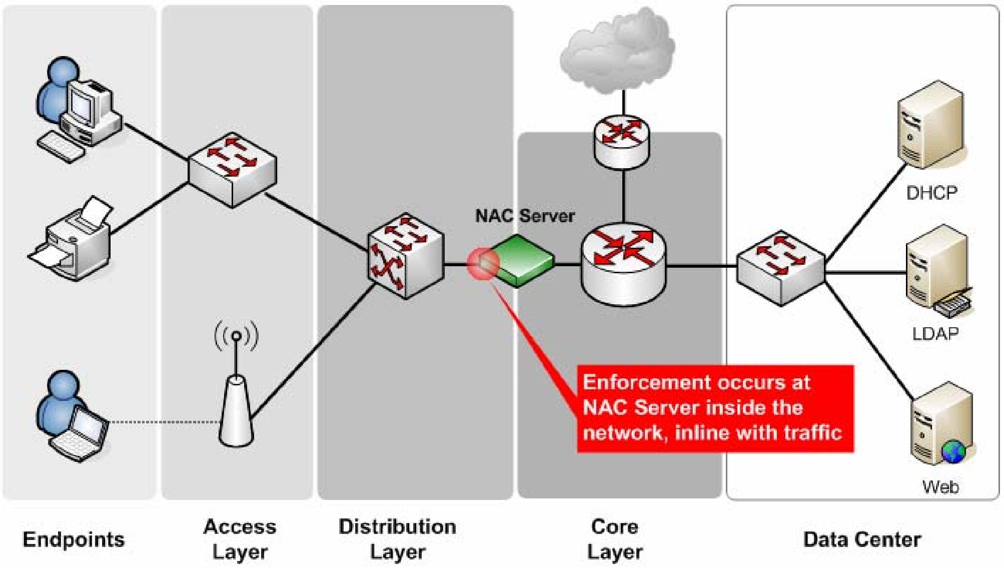

Figure 2 – In-Line Enforcement

Hybrid Enforcement

Hybrid Enforcement, common in early NAC products, combines both in-line and edge enforcement techniques. In hybrid enforcement, an end user's system starts "behind" the NAC server, which is sitting in-line, enforcing access controls itself directly. Typically, the NAC server will be controlling access, requiring authentication and validating end-point security posture. Once the authentication and posture assessment is complete, the in-line device takes itself out of the critical path between the user and the network. It does this by reaching out to the actual edge switch the user is directly connected to and modifies the configuration to apply the appropriate enforcement. For example, the NAC server might change (in the edge switch) the VLAN the user is connected to, moving it away from the in-line NAC server and onto the appropriate VLAN.

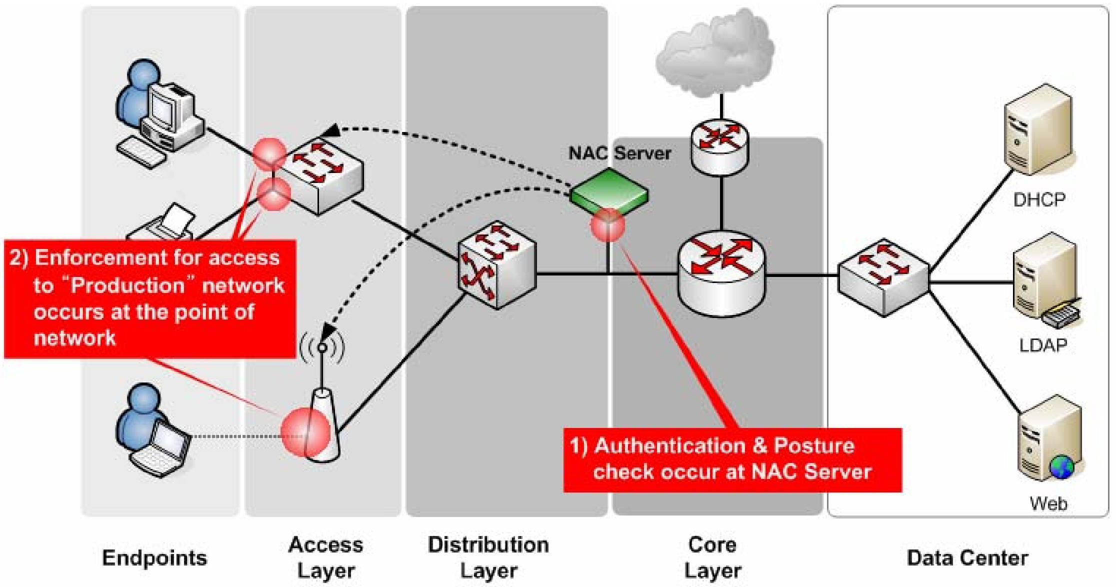

2) Enforcement for access to "Production" network occurs at the point of network

1) Authentication & Posture check occur at NAC Server

Figure 3 – Hybrid Enforcement

Protocol-based Enforcement

Protocol-based enforcement, a very low security approach, is normally used when the goal of NAC is primarily compliance checking and not enforcing access control. The most common case of protocol-based enforcement uses DHCP as an "enforcement" mechanism by giving users addresses, subnet masks, or gateways that restrict their network access in some way. For example, a user might be given an IP address on a subnet that only lets them talk to the NAC server for authentication and end-point posture assessment. When these phases are complete, the user gets another IP address that gives them access to the rest of the network. However, this is all done with protocol elements within the DHCP dialog: nothing except the willingness of the user not to cheat actually "enforces" the access controls. Other protocol-based enforcement mechanisms include deliberate manipulation of ARP caches or MAC forwarding tables to restrict access.

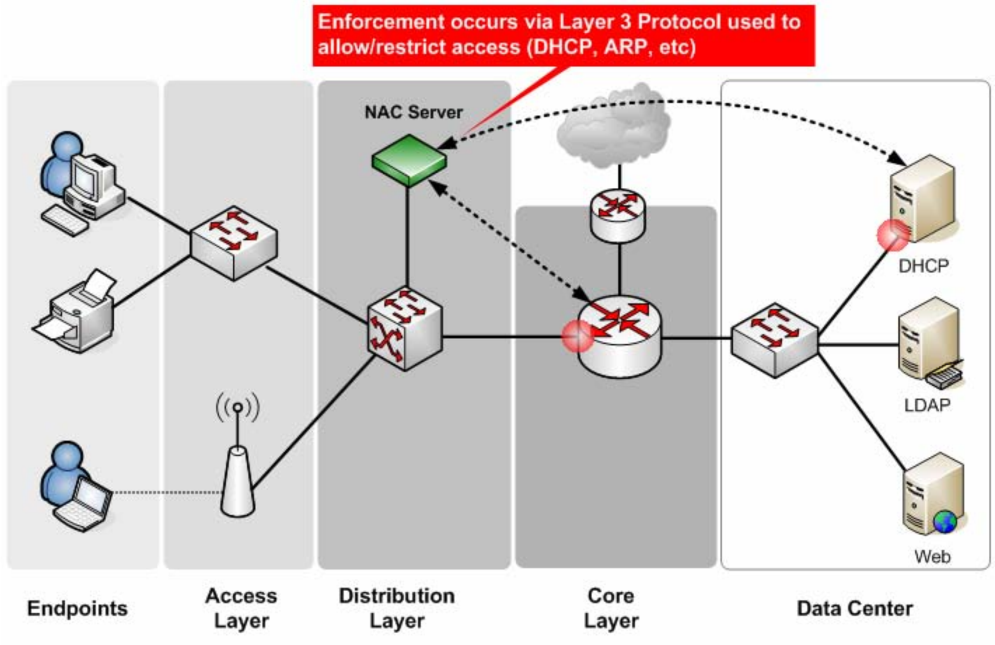

Enforcement occurs via Layer 3 Protocol used to allow/restrict access (DHCP, ARP, etc)

Figure 4 – Protocol-based Enforcement

Implementing one of these approaches does not preclude the ability to implement the others. In fact, it is common to see combinations of these approaches in real-world NAC deployments. However, each of these enforcement approaches has some specific advantages and disadvantages. While some network topologies or NAC use cases will be restricted to one approach or another, many network managers have a clean slate for their NAC deployment and need more information about the pros and cons of each approach and which is the best for their network.

In this white paper, we define five key characteristics, including security, flexibility, risk, scalability, and cost, of these different NAC enforcement techniques. By considering each characteristic separately, we can bring the strengths and weaknesses of the different approaches into clear focus.

SECURITY

NAC is primarily a security service, and thus the most important criteria for selecting a NAC enforcement method is the security of the enforcement.

It may be easy to claim that one approach or another is “more secure”, but supporting that claim can be tricky, mainly because security is difficult to measure. The only way to gauge the security of some technology is to determine how well it supports the security goals of the organization. In simpler terms, a technology is secure if it does what you want it to do.

With this in mind, how can we determine how secure these different NAC approaches are? The answer is that we find out how well they do what we want them to do which is 1) enforced authentication and, 2) the process of binding authentication to device. Let’s consider them one at a time.

Since access control is based on the user’s authentication information plus the end-point security posture of their device, clearly the user and the device should have the absolute minimum ability to see or use the network until the access control function has been fully computed. There is an obvious ordering of authentication first, and access control enforcement last.

Edge enforcement best supports the enforced authentication goal of NAC. With this approach, authentication occurs before the user actually connects to the network. Edge enforcement ensures that the user cannot send any packets using any protocol anywhere in the network until the authentication dialog has completed. The user doesn’t get an IP address; they can’t send normal packets; they can’t send “crafted” packets (such as those designed to fool switches or other devices); they can’t send anything at all. Edge enforcement blocks all access until after the user has authenticated.

Likewise, users cannot receive any packets. They can’t sniff other users’ traffic; they can’t see ARPs or multicast packets; they can’t see anything, because they aren’t on the network. Even in a wireless network using 802.11i (WPA/WPA2) encryption, the user can’t see anyone else’s packets, because everybody gets a different encryption key—and not until after the authentication dialog has successfully completed.

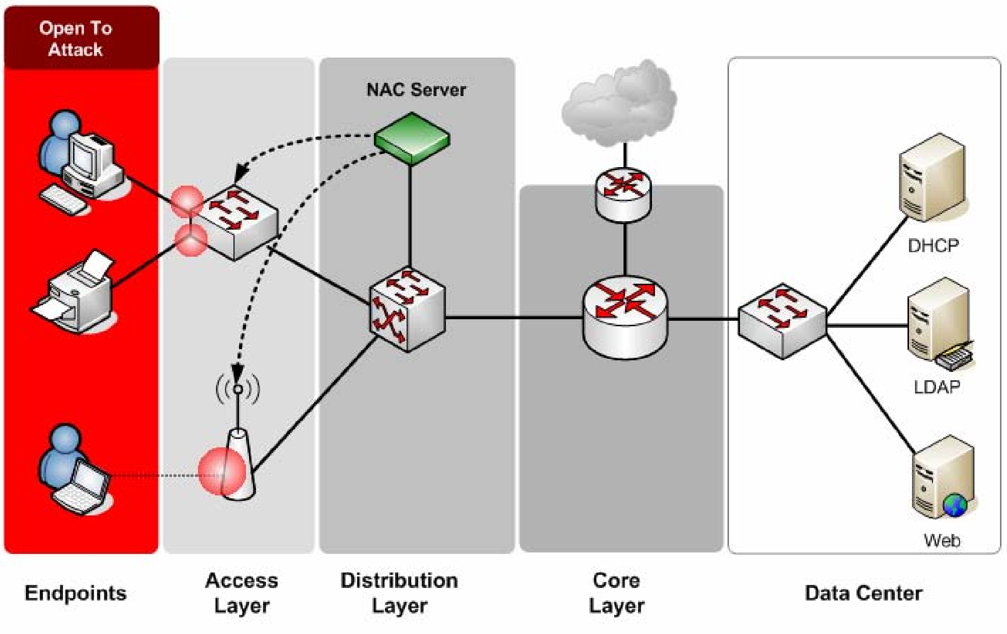

Figure 5 – Area of control for Edge Enforcement

In-line, hybrid, and protocol-based models all offer weaker security than edge enforcement because in these other three NAC scenarios, the user is fully participating in the network before they begin the authentication dialog. With 802.1X and edge enforcement, the security can be even stronger because both authentication and posture assessment are completed before the user gets onto the network. Using in-line or protocol-based models, the user gets an IP address; the user can see (possibly all) packets going through the switch; the user can even send packets (although they might not get very far). In these approaches, the user is free to attack the network infrastructure—including the NAC server itself. Of course, the user can't get past a NAC server blocking access, but they might be able to attack other users coming onto the network who have not yet authenticated or who are in a quarantine state.

Products sporting only protocol-based enforcement have the biggest hurdle to overcome as their successful enforcement primarily rests on the users' willingness to play by the rules so as to not break the overall NAC scheme.

In-line and hybrid enforcement products, however, do have a variety of techniques that attempt to minimize the ability of an unauthenticated user to wreak havoc on the network. For example, their makers harden the NAC server as best they can to minimize the possibility of a successful attack on the server itself. They isolate users as much as possible in an authentication or remediation VLAN to keep unauthenticated users away from authenticated users.

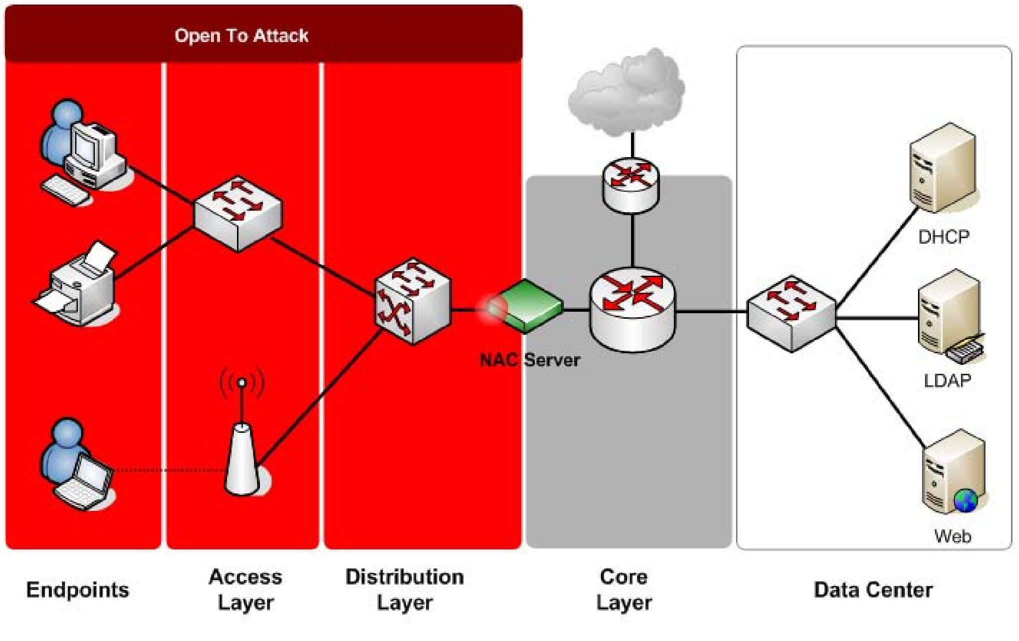

Figure 6 – Area of control for Other Approaches

There are other security differences between these approaches as well. For example, in hybrid enforcement, the network infrastructure, such as an edge switch, is only loosely bound to the NAC policy decision point and control system. When a user plugs into or unplugs from the network, there is nothing on the edge switch which immediately changes their access parameters. It's up to the NAC control system to notice these changes and re-configure the switch. The mechanisms that hybrid enforcement vendors are putting into place to help improve the reliability and security of these systems, such as constant polling of switches or constantly PINGing end users, have the undesirable side effects of increasing the load on switch management engines—never high powered to begin with—and increasing the overall network load, which can be a significant factor if NAC is used in a wireless environment.

If you rank these different approaches against the goal of providing positive and continuous control over user access to the network, it's obvious that edge enforcement has the greatest level of security, with in-line, hybrid, and protocol-based enforcement offering progressively less security.

A second aspect of security is the binding of the user's authentication to the actual network device. Many of the recent vulnerabilities discussed in the security community center on binding problems: offering an inappropriate level of trust to a user or device in the mistaken belief that you know who or what is on the other end of the line. These vulnerabilities have many different forms, but they are usually described as "man in the middle" (MITM) attacks.

802.1X edge enforcement offers the strongest binding of the user authentication to the actual connection because there is very little "middle" to have a man in. With edge enforcement on LAN switches, the device enforcing access is connected via a patch cable to the user's system. The simple truth is that the shorter the link between the user's device and the enforcement point, the smaller the window of vulnerability.

In-line and protocol-based enforcement have the weakest binding between device and presumed identity because the enforcement device sits much deeper in the network than an edge device. Hybrid authentication falls somewhere in between the two; while the authentication is loosely bound to the device, the actual enforcement happens much closer to the edge.

Consequently, this weak binding has two security vulnerabilities. One common to in-line, hybrid, and protocol-based enforcement is in the authentication itself, where a “man in the middle” could steal credentials from the user for later reuse. The other security vulnerability is more specific to in-line and protocol-based enforcement, because an attacker can insert packets into the network ahead of the enforcement point using a forged MAC address or forged IP address—all undetectable to the enforcement point.

By this measure, edge enforcement again offers the highest level of security, followed by hybrid approaches, with in-line and protocol-based trailing in their level of security.

FLEXIBILITY

Every security architect desires the maximum flexibility in the products they pick to deploy. Getting locked into “one way of doing things” is dangerous, and limits the ability to respond either slowly or quickly to new threats or new business requirements. NAC enforcement approaches that have greater flexibility are more desirable, especially with regards to evolving security technology.

The key area in NAC enforcement approaches that calls for great flexibility lies in the method by which enforcement can be accomplished. For example, one enforcement method offered on any NAC enforcement approach is “go/no-go enforcement” meaning that you either can get on the network or you can’t. At the other end of the spectrum of enforcing access, some products use full stateful firewalling as a method for enforcing access.

These enforcement methods can be used at different points in the network, so there is not a one-to-one mapping between enforcement approaches, such as edge or in-line, and the method of enforcement, such as “go/no-go” or firewalling. However, not every enforcement approach works with every enforcement method.

Edge enforcement approaches offer the greatest range of enforcement methods, and thus the greatest flexibility, by allowing the security architect to pick from a spectrum of enforcement choices. Depending on device capabilities, edge enforcement can include a “go/no-go” decision on network access, VLAN assignment and segmentation, simple packet filters, and full stateful firewalling. The most common edge device in enterprise networks, the Cisco Catalyst switches, offer three of these four options immediately out of the box, stopping only short of full stateful firewalling—although even this option is available to some extent when the switch is upgraded to include Cisco’s firewall service module.

Hybrid NAC schemes have a different type of flexibility. Instead of being flexible in the method of enforcement, they are more flexible in the location of enforcement. Because the hybrid NAC approach depends on changing the configuration of network infrastructure, they offer the flexibility of choosing where in the network you want to have enforcement occur. However, just because hybrid NAC approaches can offer flexibility in location, doesn’t mean that every existing NAC product has that level of flexibility.

Network architects who want maximum flexibility in the method of enforcement, such as VLANs or firewalling, should concentrate on edge enforcement schemes. Those who want maximum flexibility in the location of enforcement should look at hybrid approaches. And, of course, for even greatest flexibility, some vendors combine these approaches by enforcing different types of access at different points within the network.

In-line enforcement by itself offers the security architect less flexibility because the device enforcing access is deeper in the network and cannot offer all these options. For example, in-line enforcement cannot completely block a user from all network access(go/no-go enforcement), and cannot provide full VLAN isolation of users (since the traffic must be trunked to the enforcement device on a common

VLAN for all users). This lack of flexibility in in-line devices is very common, largely because in-line devices are aimed at simpler, smaller deployment scenarios where “flexibility” is not an important criterion for the overall NAC solution.

The least flexible approaches to NAC enforcement are protocol-based ones. Because these approaches are based on the particular behavior of a protocol (such as the sequence of messages in a DHCP address request or renewal), they usually only have go/no-go access controls.

Another area of flexibility worth considering in NAC enforcement approaches is the ability to control more than IPv4-based traffic. Although it is common for network managers to think exclusively in terms of IPv4, TCP, and UDP, enterprise networks have a much broader range of protocols running across them—intentionally or unintentionally. Certainly, network architects are beginning to look forward to IPv6 and current versions of Windows desktop, Macintosh, and Linux operating systems come equipped out-of-the-box with IPv6 support enabled. When Longhorn (Windows Server 2008) ships, it will also have IPv6 capabilities from the moment it is installed. With edge-based enforcement, the authentication and posture assessment decisions are all made independently of the protocol that will later be used, whether it’s IPv4, IPv6, or something else that might come down the pike.

Some enforcement approaches, especially in-line and hybrid, don’t necessarily eliminate the possibility of using non-IPv4 protocols. However, the flexibility of edge enforcement becomes obvious here: edge enforcement naturally supports any protocol you’d like to run, while in-line and hybrid enforcement will only include IPv6 support (for example) when sufficient customers scream loud enough that it gets moved to the product.

The natural flexibility of edge enforcement also plays into the need for enterprise networks to have support for both legacy and future protocols. Most large networks have a wide variety of protocols floating around on them, ranging from Novell IPX to DECnet and SNA. Edge enforcement naturally controls all of these protocols and access methods, while in-line enforcement can only cover the protocols that are explicitly supported in the product. This enforcement is limited to IPv4 controls in current products—often ignoring the non-IP holes that hackers love, such as ARP.

PRAGMATIC DEPLOYMENT AND RISK OF FAILURE

Taking a pragmatic, step-by-step approach to deploying NAC really increases the chance of success. It is better to deploy a NAC solution in small pieces, merging it cleanly with the network, increasing the level of security and control as you grow comfortable with the reliability and stability of the solution. The alternative, going for a watershed day cutover of large parts of the network, has considerable risk of failure with anything as complex as NAC.

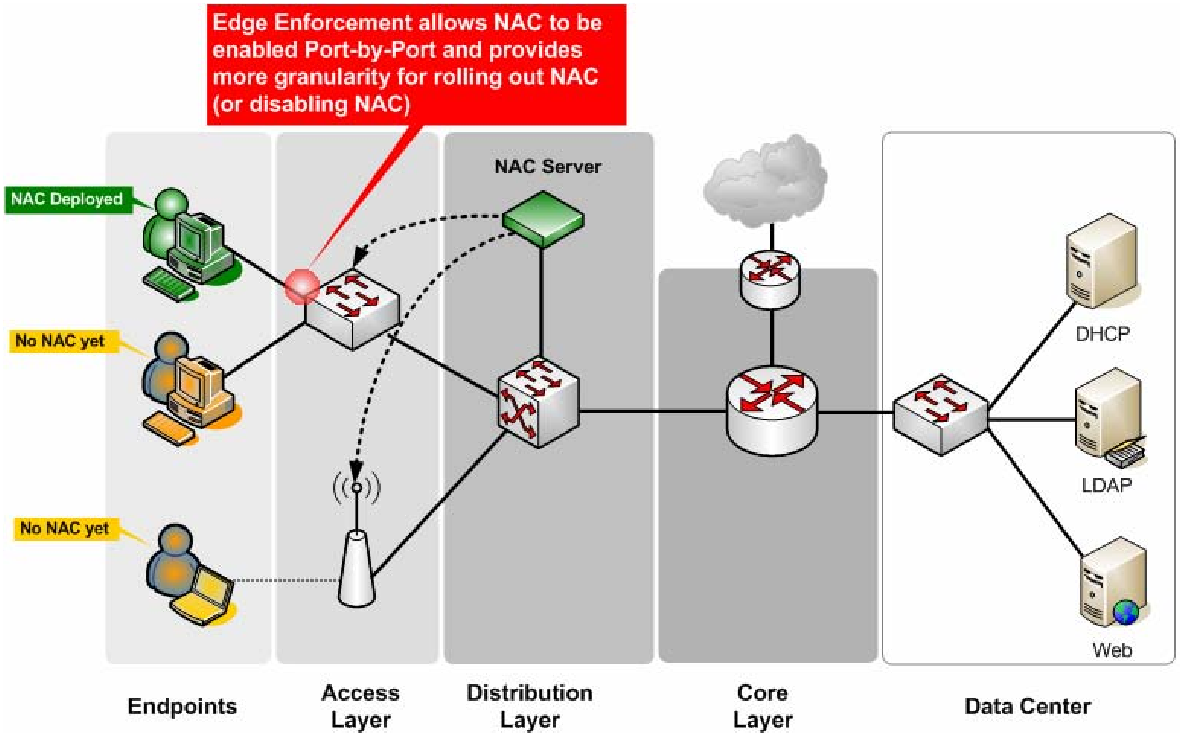
Edge Enforcement allows NAC to be enabled Port-by-Port and provides more granularity for rolling out NAC (or disabling NAC)
NAC Deployed
No NAC yet
No NAC yet
NAC Server
DHCP
LDAP
Web
Endpoints
Access Layer
Distribution Layer
Core Layer
Data Center

Figure 7 – Deployment Flexibility of Edge Enforcement

If step-by-step refinement and a gradual deployment are important to the success of a NAC solution, then choosing an enforcement approach that supports this is crucial.

Because edge enforcement maps directly to network devices already in place and under management, it inevitably brings a smoother deployment and has the greatest chance of success. When the authentication and NAC features can be enabled on a port-by-port basis, switch-by-switch, the NAC architect has the greatest granularity to bring NAC services into play (and out of play) at the point where the users connect to the network.

While edge enforcement has the advantage of port-by-port deployment, there are scenarios where very simple topologies using in-line installation will be, in fact, simpler. For example, if NAC is primarily aimed at guest users who might be sitting in conference rooms all plugged into a specific switch or everyone coming into the network via wireless links, the simplicity of an in-line installation is fairly obvious. A fully encapsulated in-line solution is easier to deploy and less likely to cause disruption than either edge or hybrid enforcement, both of which require edge switch changes.

One advantage that edge enforcement has over in-line enforcement, though, is the ability to quickly and easily back out of network changes. While no one ever wants to take a step backwards, having a plan to "undo" a change to the network if there are problems is a common-sense requirement. Any NAC solution will require the addition of policy servers to a network. Once these are installed, activating (or deactivating as the case may be) edge enforcement is easy, because turning NAC on and off on an individual port is as simple as issuing a single command line or ticking a check box in a GUI.

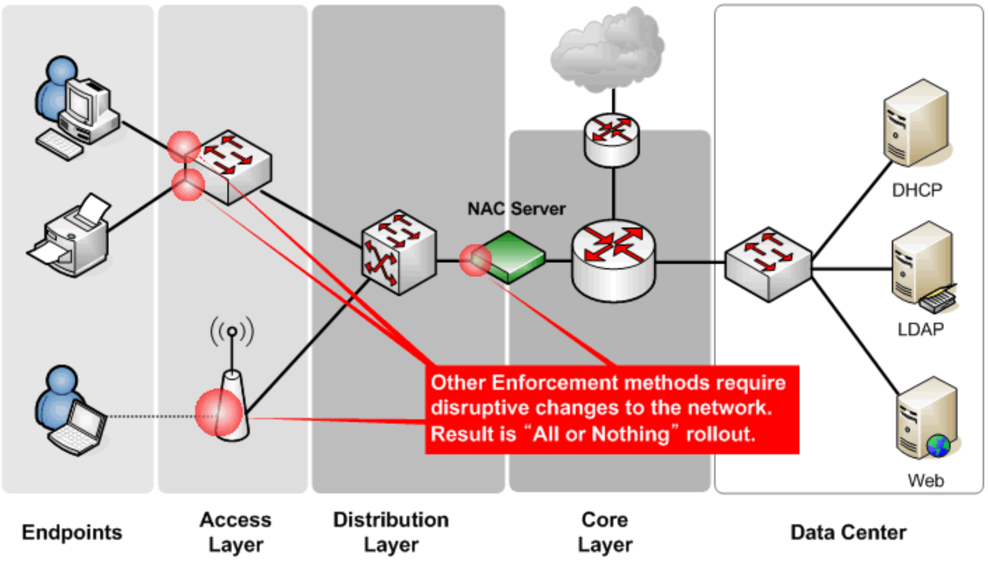
NAC Server
DHCP
LDAP
Web
Other Enforcement methods require disruptive changes to the network. Result is "All or Nothing" rollout.
Endpoints
Access Layer
Distribution Layer
Core Layer
Data Center

Figure 8 – Deployment Flexibility of Other Enforcement

In contrast, in-line enforcement causes a network topology change: the in-line device, after all, has to be placed in-line. As an intrusive technology, there is no way to install on in-line device without affecting large numbers of network users. In-line enforcement represents a greater commitment towards the NAC solution. When an enterprise LAN is misbehaving, you may not even get the chance to examine the issue, do the necessary debugging and subsequently fix the problem. If it doesn't work, the only thing you can do is pull it out.

Hybrid enforcement, unfortunately, combines all of the risks of both edge and in-line enforcement. If the in-line part of the system doesn't work, no other NAC function works either. For configuration complexity, hybrid is even more difficult to set up than edge enforcement, because it requires an intimate linkage between the NAC solution and the entire topology of the network. Hybrid enforcement solutions, because they reconfigure switches on-the-fly, represent a technology which maps least cleanly to how networks are managed and configured, offer the lowest granularity, require the greatest commitment, and entail the most configuration complexity.

Protocol-based enforcement methods, because they vary significantly in their technologies, are difficult to place in this spectrum. If we use the example of a DHCP-based approach, then the changes to the network occur at multiple layers,

SIMPLE... BUT WHERE?

There is considerable debate, at least among NAC vendors, about the relative simplicity of different authentication approaches. Edge enforcement is normally associated with 802.1X supplicants (although other techniques, such as MAC address-based authentication, are also common), and the relative lack of experience of network managers with 802.1X supplicants makes them an easy target for FUD attacks.

The alternative to 802.1X supplicants proposed by NAC vendors is often called "dissolving client" or "zero touch", meaning that the end user's system isn't touched by the network manager to prepare it for the NAC deployment. The burden of touching, in this case, has been moved to the network management and configuration side. For example, non-802.1X solutions typically need an in-line component (whether hybrid or in-line), as well as management links to user switches so that they can see link transition events, read MAC and ARP tables, and reconfigure switches and switch ports as users come and go from the network.

Many "zero touch" solutions also require access to network taps or mirror ports so that they can inspect all traffic passing by and sniff Windows login events as well as other relevant traffic. They may also need hooks into Windows domains so that they can map identity information to a connection.

The preference for one approach or the other will depend on the relative pain and disruption that each may cause. NAC authentication using 802.1X or a proprietary client disrupts end-users, and minimizes disruption of the network. Authentication using other approaches moves the disruption from the desktop team to the network team. While anything but "zero touch", these solutions have found their place in the market in environments where substantial network changes are less disruptive than installing a client or using 802.1X.

both in the infrastructure itself (by having multiple subnets running on a single network segment) and in the DHCP service on top of that infrastructure. As with an in-line deployment, many of these changes are an all-or-nothing update: if things don't work, then there is little alternative but to put everything back to the pre-NAC days. Debugging these types of changes will also be difficult in all but the simplest network topologies.

DISTRIBUTED ENFORCEMENT, PERFORMANCE, AND SCALABILITY

Two potential dangers zones for any security technology are scalability and false positives. NAC technologies that help to minimize these dangers will be easier to deploy and more successful post deployment. It is important to remember that the goal of NAC is to get people onto the network, not to keep them off it. This implies that the NAC solution should be robust and reliable. If NAC stops working, so does everyone in the company.

Edge enforcement, and to some extent hybrid enforcement, benefit greatly from the distributed nature of the enforcement mechanism. Enterprise-class switches are engineered for a particular load based on a number of ports and traffic levels. They don't represent a bottleneck when operating within the engineered limits. In-line devices, especially those based on general purpose computers, can easily become a bottleneck. It's safe to say that a switch doing VLAN enforcement will never cause a performance problem in a NAC deployment. It's less obvious that in-line devices, especially those based on off-the-shelf servers or those working with processor-intensive advanced security features, can handle the load in a cost-effective manner.

There are scalability issues with edge and hybrid enforcement which are related to the load on the switch's CPUs. Most switches are designed with very inexpensive control systems and very fast switching systems. This is because common control operations, such as configuration changes, happen weekly or monthly, and there's no point in spending a lot of money putting in a fast control system for something that doesn't happen very often. In edge enforcement, some additional load is placed on the switch when an authentication occurs. Hybrid enforcement can put a much higher load on the switch by using tools such as SSH/Telnet and SNMP to closely synchronize the state of the switch with the NAC policy servers. Depending on the design of the hybrid NAC system and the age of the switches, this additional load can cause uncharacteristic performance problems.

There are always going to be scalability challenges in any NAC deployment. In in-line enforcement models, the challenges are very obvious because of the in-line nature of the enforcement. Hybrid systems have their own scalability challenges. They are in-line during part of the connection, and also have a higher performance load because of their constant communications with the distributed switches. Hybrid systems that also monitor traffic can suffer severe scalability problems because they must monitor large volumes of traffic. Monitoring represents a challenge because it is often difficult to find places to monitor that traffic that are not already occupied by IDSes or other security tools. While hybrid enforcement approaches vary in their architecture and design, they will always be less scalable than simple edge-based enforcement, sometimes dramatically so.

In edge enforcement, while there are still potential bottlenecks (such as where the policy and authentication servers are concerned), the load of enforcement has been spread very cleanly across the network switching fabric. This reduces the likelihood of scalability and performance problems.

Protocol-based enforcement methods can also be severely restricted in their scalability. For example, some protocol-based methods depend on timing of packets (such as those based on controlling ARP caches and MAC forwarding tables) and tend to burst large numbers of packets at very high speeds throughout the network to achieve their enforcement goals, which can adversely affect switch performance. Similarly, when the packets are multicast, they can have huge performance effects, particularly on large distributed LANs, where every switch and every device in the network will see and have to act on the packets. However, because protocol-based enforcement is normally aimed at small networks (such as SMBs or branch offices), these scalability issues may be more potential than actual.

False positives are a more subtle problem in NAC deployments, but it is still one for which you need to be aware. Any security test is going to yield false positives some percentage of the time. As security architects, it is our job to work to minimize them though. With edge-based enforcement, the NAC designer achieves an especially trustable knowledge of a particular switch port (layer 1), MAC address (layer 2), and IP address (layer 2-3 binding). When such confident knowledge is available, the system architect can have a higher confidence in different security controls. This, by itself, can give a lower level of false positives.

In addition, because the knowledge that a particular IP address is bound to three things: a particular MAC, user, and port, is known with high confidence, this information can be reused as part of other security policies. For example, some edge-based enforcement strategies are combined with deeper enforcement processed at firewalls placed in data centers and in front of sensitive servers located around the network. When you really know who is behind a particular IP address, you can more safely install security policy on devices distant from the actual end user.

With in-line and hybrid models, the level of confidence in topology and binding is correspondingly lower. For example, a switch will know instantly when a MAC address moves from one port to another—a potential sign of some security funny business. But on the other hand, an in-line or hybrid solution will have to either rely on getting notification from the switch or depend on CPU-stressing polling operations adversely impact the performance of the switching infrastructure. Because the confidence level is lower, the chance for error, undetected attacks, and false positives is correspondingly higher.

Hybrid and protocol-based enforcement models are particularly vulnerable to deliberate deception on the part of network attackers. Because they depend on unreliable “best effort” packet delivery and a wide variety of clever management and detection tricks, they simply escalate the war with a more clever trickster in the attacking foe. When a security enforcement system is not fundamentally built into the network, but layered on top of an existing infrastructure, there will be greater susceptibility to network changes and determined attackers.

COST

A good strategy for NAC in the enterprise is to leverage of value in the equipment that is already deployed. There are two ways this can, and should, occur. One is to simply use the features that are already present in installed equipment, such as 802.1X authentication and VLAN capabilities, both present in all enterprise switches as far back as 2002. The other is to build on the strength of network infrastructure vendors, all actively looking to improve the security of their products.

Working with what you already have, already know, and already trust represents a great cost savings. Switches and existing network topologies generally already support edge enforcement via 802.1X, both in wired and wireless environment. In fact, most secure wireless networks already make use of 802.1X. VLAN-based enforcement also depends on a feature that enterprise vendors have been including in their equipment for many years. Because the equipment with these capabilities is already installed and familiar to the network operations team, this reduces the potential capital expenditure associated with an edge-enforcement deployment.

A second cost savings is to leverage features in your existing switching infrastructure specifically aimed at security. Network equipment vendors are constantly getting “smarter” about security, thus their switches and routers are bringing new, high-end security features to the table. A NAC strategy that leverages these advances in hardware and software is especially cost-effective—and edge and hybrid enforcement naturally align. Standards-based NAC using edge or hybrid enforcement builds on the security expertise and products of the entire industry, rather than depending on the products and services of a single security vendor.

Operational costs are also lower with edge enforcement. With in-line enforcement, invisible NAC devices act as firewalls deep inside the network. The very traditional and well-known tiered hierarchy used in building large LANs is broken up into islands of enforcement, with new NAC devices in between LAN switches. These newly inserted NAC devices are not the same type (or even vendor) of device, which means that normal network management tools which capture logging information, display topology, and handle the network as a whole are suddenly incapable of seeing all the critical network elements. You can't debug a network problem anymore by looking at switches; now you also have to talk to in-line enforcement firewalls and NAC policy servers to hope to figure out what's happening in the grander scheme of access control. Inconsistencies in the network will cause continuing operational costs.

Compare this to edge enforcement, where every aspect of a user's port configuration is located on the switch itself, easily visible using existing operations and management tools. A standardized approach will, over the long run, reduce operational expenses.

In-line enforcement always requires additional hardware; it can require hardware of nearly equivalent cost and complexity to the edge switching infrastructure it supplements. Even in hybrid models, which eventually use edge enforcement, additional costs of in-line authentication and posture checking devices are required.

What in-line NAC solution vendors don't like pointed out is that in-line approaches simply cost more money. Why spend that extra money when many enterprises can build a powerful NAC solution by adding nothing more than a NAC policy server to their existing network.

Protocol-based enforcement methods tend to minimize the amount of additional hardware they require, being roughly in the same order of magnitude as hybrid enforcement systems. However, cost-wise, they are less effective because they don't really leverage the hardware you've already bought—they come up with ways to work around what you have already paid for and deployed.

CONCLUSIONS

A wise security architect looks to provide high security, good flexibility, low risk, predictable scalability, and reasonable cost in the solutions they design. While every deployment scenario is different and enterprise requirements can vary tremendously, edge-based enforcement an excellent choice for most enterprises looking to add NAC into their existing networks.

	Edge Enforcement	Inline Enforcement	Hybrid Enforcement	Protocol-based Enforcement
Security	Greatest level of security; enforcement occurs at the point of network access	Progressively less security; enforcement occurs deeper in the network, leaving more areas vulnerable / uncontrolled		
Flexibility	Greatest level of flexibility in enforcement methods; protocol-independent (IPv4, IPv6, etc.)	Progressively less flexibility in enforcement methods; may be dependent on behavior of particular protocol (e.g. dependence on DHCP, or may not natively support IPv4 to IPv6 migration, etc.)		
Risk	Least intrusive; with granular deployment options for lowest risk of network disruption	Changes to network topology and/or protocols are more intrusive with limited granularity (often "all-or-nothing") which increases risk of disruption to network		
Scalability	Most scalable; load of enforcement is spread across network fabric for greatest scalability and performance	Inline nature of enforcement reduces scalability and has significant impact on performance		Protocols relying on broadcasts or multicasts limit scalability and performance
Cost	Most cost-effective approach; leverages security functions of existing infrastructure to reduce capital and operational costs	Inline enforcement approach has highest capital cost (more NAC servers/appliances); operational costs are higher, particularly for troubleshooting		Similar to Hybrid capital cost; less cost-effective than Edge since existing network infrastructure is not leveraged

ABOUT OPUS ONE®

Opus One® is a consulting and information technology firm based in Tucson, AZ. Founded in 1989, Opus One's corporate goal is to help our clients make the best use of information technology. We focus on efficient and effective solutions in the areas of data networking, electronic mail, and security. For more information, see http://opus1.com or contact us at:

Opus One
1404 East Lind Road
Tucson, AZ 85719
+1-520-324-0494