Strengthen Data Protection with Network Access Controls

Tom Olzak
May 2006

In today's business information enterprise, a strong perimeter defense is not enough. Network access requirements for remote, wireless, and physically connected unmanaged systems result in solutions that bypass perimeter controls. New methods of preventing unauthorized access and malware infections are required.

As we'll examine in this paper, the first step in meeting the challenges associated with evolving demands for new network access methods is the segmentation of the network. At a minimum, network segmentation should result in a production segment and a restricted access segment. To help you accomplish this, we take a look at how VLAN's work.

Once your network is segmented, the next step is to ensure that all endpoint devices (workstations, laptops, handheld devices, etc.) conform to your security baselines. Using a generic view of network access controls, we step through three common network connectivity scenarios.

But first, let's review the reason for rethinking network security strategies—deperimeterization.

Deperimeterization

The traditional high wall of network perimeters has been breached. As depicted in Figure 1, the three primary causes of perimeter breaches are remote and wireless devices as well as generally accessible open network ports (represented by conference rooms). Blasting through the perimeter, these access paths are critical for businesses that want to effectively compete in any industry. But why should we care? Any responsible business ensures only protected systems connect to its network. OK, but this goal might be harder than you think.

Some problems with endpoint device protection

Many organizations deploy endpoint applications or other solutions that require the user to have sufficient access to laptops and other mobile devices to provide local control over security processes. Users who don't want to wait for a patch or an anti-virus update might interrupt update processes. Worse, users might decide to shut down the patch management or anti-virus software due to perceived performance issues. And even if user access isn't an issue, many remote users don't regularly connect to the company network while they regularly connect to hotel or airport networks for access to the Internet. Finally, users might download and install all types of goodies that could have

an interesting effect on your network—and these are the systems over which you have some control.

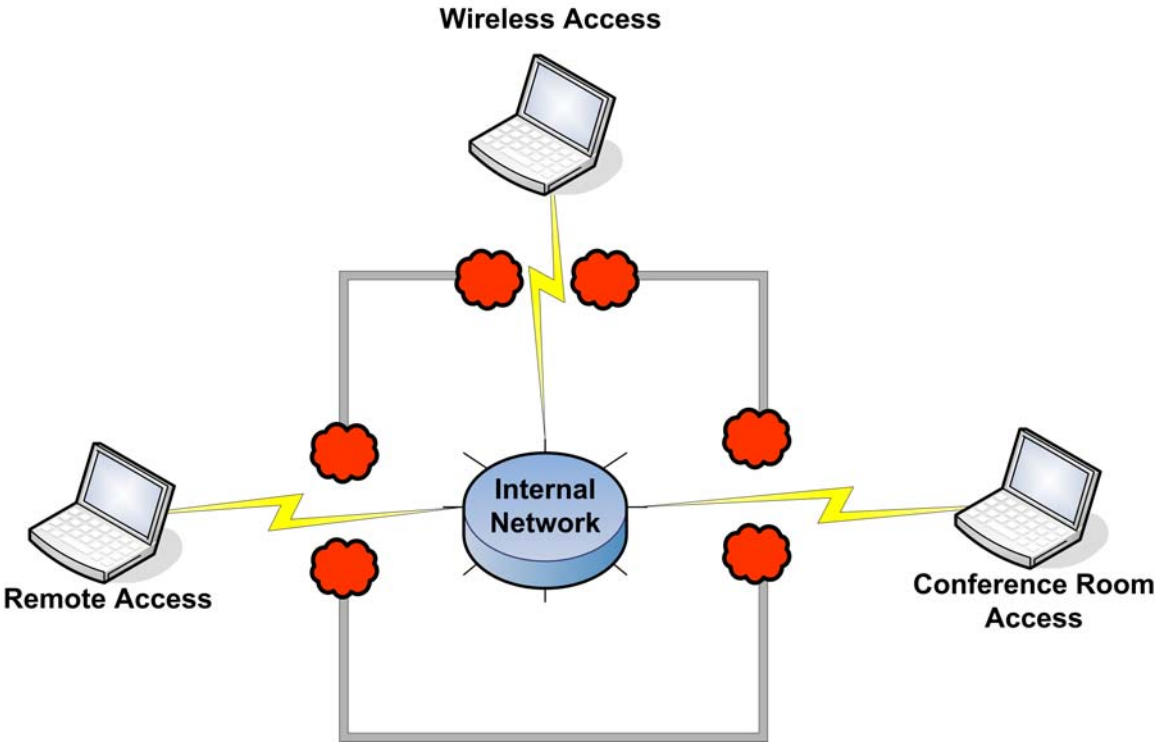

Figure 1: Deperimeterization

Vendor or customer systems are rarely managed by your security or customer support team. Risks similar to those described above might exist on non-company owned systems, but they are exacerbated by the possible absence of any protection against endpoint device compromise.

Now comes the scary part. Many organizations allow these questionable devices to bypass perimeter security controls and connect directly to the internal network. For the purpose of this paper, we're going to focus on three common network points of entry: SSL VPN, wireless, and conference room network jacks.

SSL VPN

SSL VPN access is a moderate risk. If configured properly, it becomes part of your perimeter defense. But the risk still exists that an infected system will connect to your network, passing malware through your defenses.

Wireless

Wireless access is a common way for vendors or traveling employees to gain access when at one or more of your locations. The questionable devices described earlier are also the primary users of your wireless network. Traveling employees and vendor representatives with laptops as well as employees with personal wireless devices can connect to your internal network via wireless AP's (Access Points). Connections through poorly protected AP's bypass most if not all perimeter defenses.

Conference rooms

Although I focus on conference rooms in this paper, the challenges related to general access to network jacks anywhere in your organization are identical. As with wireless, endpoint devices over which you have little or no control may connect to your network. But there is an additional issue related to open physical access ports in conference areas.

Attackers can install rogue access points, sniffers, or other devices with the intent to compromise your network. By the time someone realizes that an unauthorized device is connected, large amounts of information is probably already in the hands of someone with plans to do harm to your company.

Overall risks associated with deperimeterization

The rest of this paper describes how to protect your information assets from attacks against your network using remote access, wireless, and conference room access points. Before we move on to how NAC fits in to an overall network defense plan, let's take one more look at deperimeterization risks (Nicolett, Pescatore, and Gerard, 2004):

- A company laptop that is occasionally connected to the internal network has not been accessible for patch installation, spends time outside of the corporate firewall, is corrupted with a worm and injects the worm into the internal network.
- A home PC that is occasionally used by an employee for work purposes is corrupted with a worm and establishes a connection to the internal network.
- A laptop or other mobile platform owned by a contractor or other external entity is infected and injects the worm into the internal network.
- An employee's personal mobile device is infected with a worm and is connected to the network.
- Attackers with physical or wireless access to your network might collect information necessary to compromise your information assets.

Now that we've examined the risks, let's move to a discussion of network segmentation with VLAN's.

Virtual Local Area Networks (VLAN's)

Placing all your devices on a single network segment can cause performance issues. The first is a high rate of collisions. Collisions are caused by two or more devices attempting to user an Ethernet connection at the same time. As the collision rate climbs, network performance declines. The most common solution for this problem is the use of switches.

A switch, which is a Layer 2 device, reduces collisions by controlling which of its ports sees a specific packet. In this way, only the network segments involved in a device-to-device conversation see a specific directed packet. This significantly reduces the number of packets on any one segment resulting in fewer collisions. Many switches can also perform another function—the creation of broadcast domains.

Although implementing a switch provides for controlling directed packets, broadcast packets are still sent to all switch ports. In other words, every device connected to the switch is in the same broadcast domain. If a broadcast doesn't need to be seen by all network devices, additional segment traffic is present that unnecessarily affects performance. If your switch supports it, you can create VLAN's to segment your network into multiple broadcast domains. A simple segmented network is depicted in Figure 2.

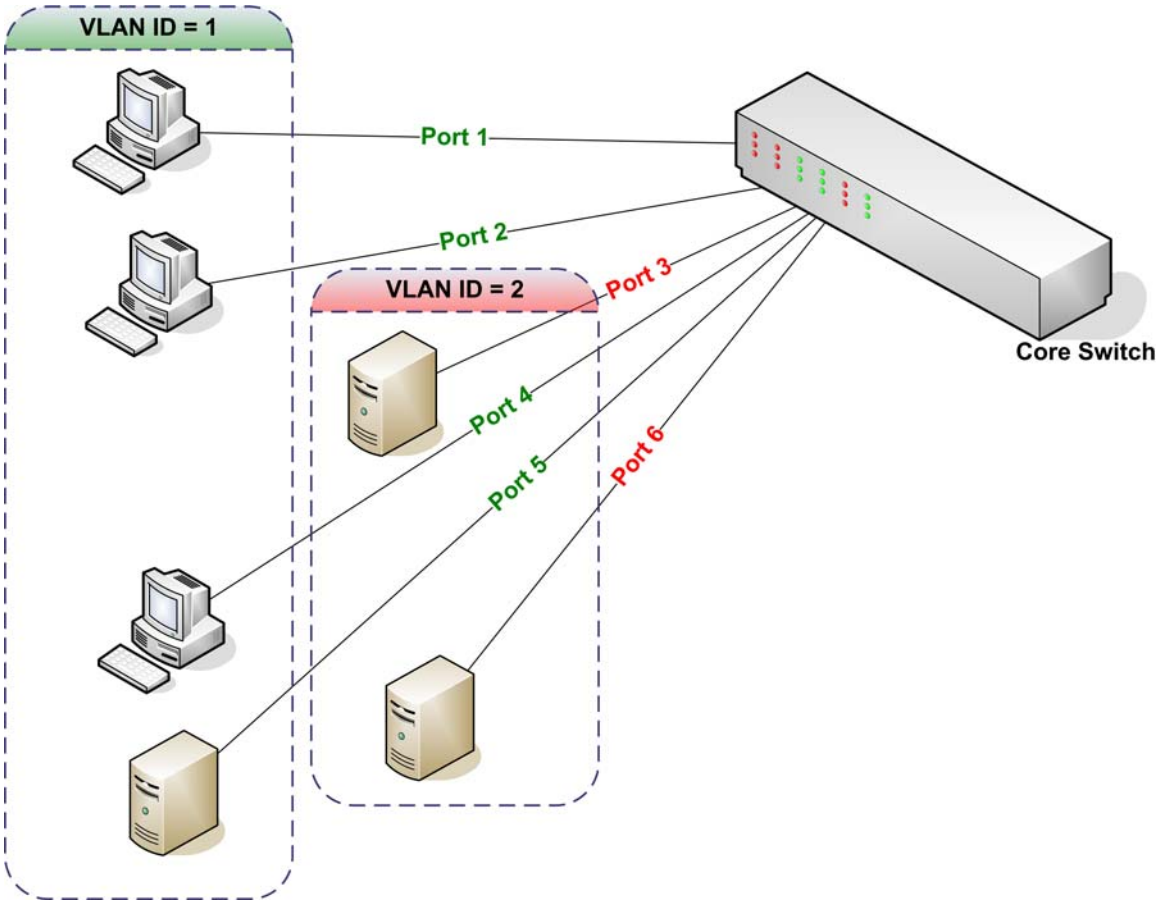

Figure 2: Simple VLAN

VLAN's are created by assigning one or more ports to a specific segment. Each segment is identified by a VLAN ID. In this example, the network is divided into two VLAN's. Switch ports 1, 2, 4, and 5 are assigned to VLAN 1. Ports 3 and 6 are assigned to VLAN 2. Neither directed nor broadcast packets placed on one VLAN are seen by the other VLAN. Besides improving network performance, there are some security advantages to deploying VLAN's, including:

- Enforcement of increasingly complex network traffic restrictions
- Implementation of logical workgroups with geographically separated endpoint devices

- Restriction of existing or new connections based on business rules provided and managed by network access control systems

Restricting packets to a specific VLAN is accomplished by using VLAN tagging. Carrying both a priority—from 1 to 7—and an ID, a VLAN tag consists of four bytes. Although VLAN-enabled stations can apply explicit tags, it's more common for implicit tags to be attached to a packet by a switch (Phifer, 2002).

A switch knows the proper ID for a packet based on the port through which it receives the packet. So in Figure 2, any packet entering the switch via Port 1 receives a VLAN ID of 1. Because of the implicit tag applied by the switch, the packet can only communicate with network resources on VLAN 1. We'll cover moving packets from one VLAN to another later in this paper.

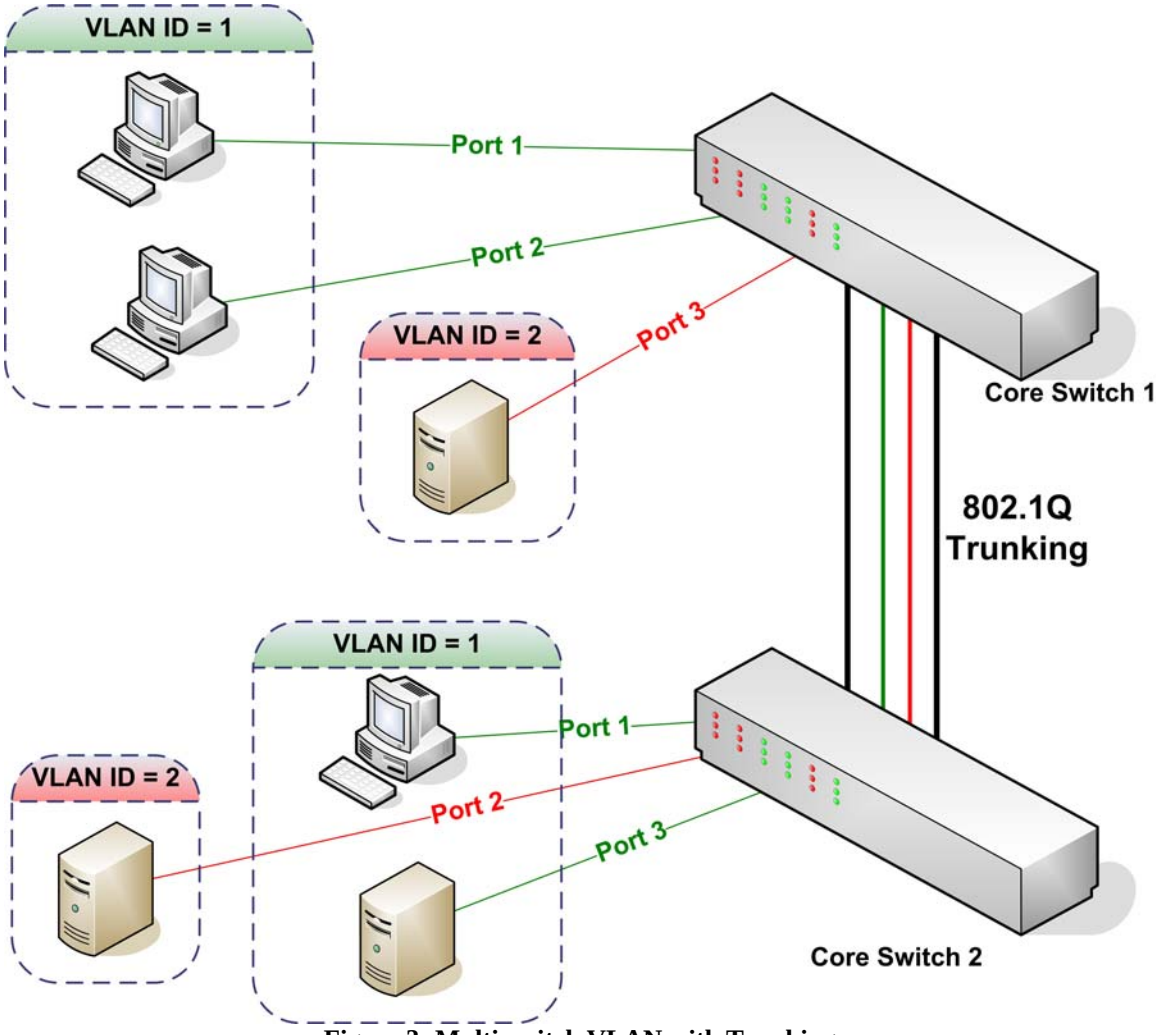

Figure 3: Multi-switch VLAN with Trunking

Up to this point, we've looked at VLANs connected to a single switch. Figure 3 shows a multi-switch configuration. Using this configuration provides significant flexibility. For example, the workstations in VLAN 1 on Core Switch 1 could be located

in the Engineering building on a university campus, while the workstation and server connected to Core Switch 2 are used in the administrative building across campus. This is called a distributed VLAN.

VLAN traffic passing between the two switches is managed by the IEEE 802.1Q protocol. 802.1Q Trunking allows packets to flow between switches with VLANs in common.

Using VLAN's for Restricted Network Segments

VLAN's are ideal for restricting traffic. For example, a logical DMZ can be created by using a VLAN. Packets entering the DMZ from the Internet are assigned a restricted VLAN ID that allows access only to devices on the DMZ. So far this is no different from a standard physical DMZ. The difference is the flexibility of the VLAN approach. DMZ devices on a VLAN don't have to be physically located together. DMZ servers can be located anywhere on the enterprise network—as long as they all share the same VLAN ID.

Carrying VLAN's a step further, restricted network segments can be created for workstations that are identified as infected with malware or that fail to meet specific security requirements. Because of the distributed nature of VLAN's, endpoint devices from anywhere on the enterprise network can reside on a single restricted segment.

Routing VLAN's

So far, our VLAN examples have been configured on Layer 2 devices. So the packets from one VLAN could not pass to other VLANS. To route VLAN tagged packets between VLAN's requires standard IP routing. An example of a routed VLAN network is shown in Figure 4.

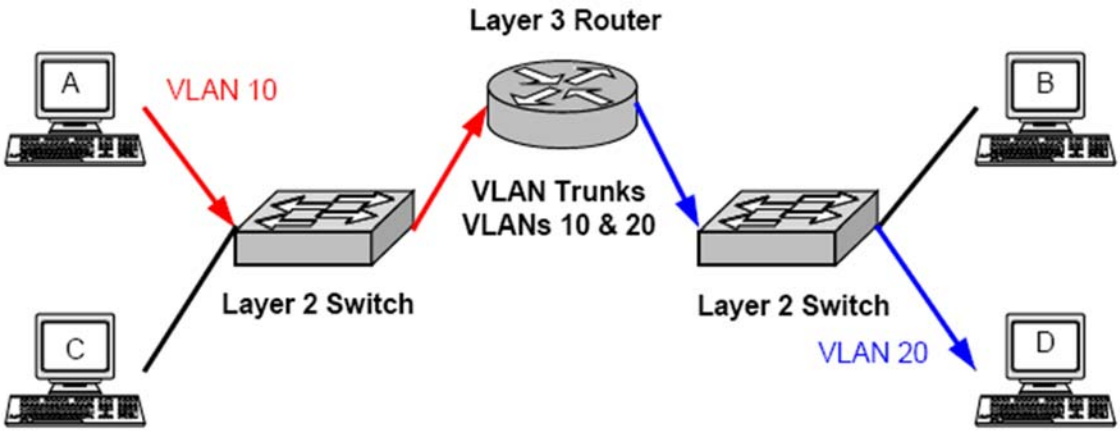

Figure 4: VLAN Routing (Dell, 2004)

Each VLAN represents a subnet with an IP Address range unique to the organization's network. When a packet from VLAN 10 is addressed to an endpoint device in VLAN 20, it's routed using the target IP address. Either a Layer 3 switch or a router can be used.

VLAN technology provides the ability to easily restrict or quarantine access by non-compliant or infected endpoint devices. Now we'll look at how to identify and dynamically configure these devices.

Network Access Control (NAC) Basics

According to Gartner, there are two business objectives when deploying and managing NAC solutions (Nicolet, Pescatore, and Gerard, 2004):

1. Only allow connections by endpoint devices that are judged to be safe
2. Detect the actions of a malicious device and quickly isolate it from the rest of the network

In other words, don't allow high risk devices access to your primary production network segments, and remove high risk devices that are already connected. In this paper, we look at two NAC approaches—scan/block and scan/quarantine.

In a scan and block scenario, high risk devices are simply denied access to any network segment. Messages might be sent to the user during access attempts notifying him about what he needs to do to lower the risk level on his device, but no connection is made that would allow remediation using the organization's network.

When an organization uses a scan and quarantine approach, high risk endpoint devices aren't completely blocked from the network. Instead, they're connected to a restricted segment where the user can either manually update her system with current patches, anti-malware software, etc., or the device uses an agent to automatically pull updates from one or more servers connected to the quarantine network segment.

Planning for and implementing a NAC solution requires planning, design, deployment, and management activities. These activities are the subject of the next section.

NAC Implementation and Management Process

NAC planning and management can be complex, depending on the organization. Figure 5 shows the NAC Implementation and Management Process (NACIMP) with phases in which specific tasks are performed and deliverables created. The objective of this process is the creation and maintenance of a layered NAC solution that's easy to manage and that doesn't degrade the business user computing experience. The phases in the NACIMP include:

1. Policy Development
2. Creation of Baselines
3. Access Control Design and Implementation
4. Vulnerability Mitigation
5. Network Monitoring
6. Containment
7. Maintenance

Let's look at what happens in each phase.

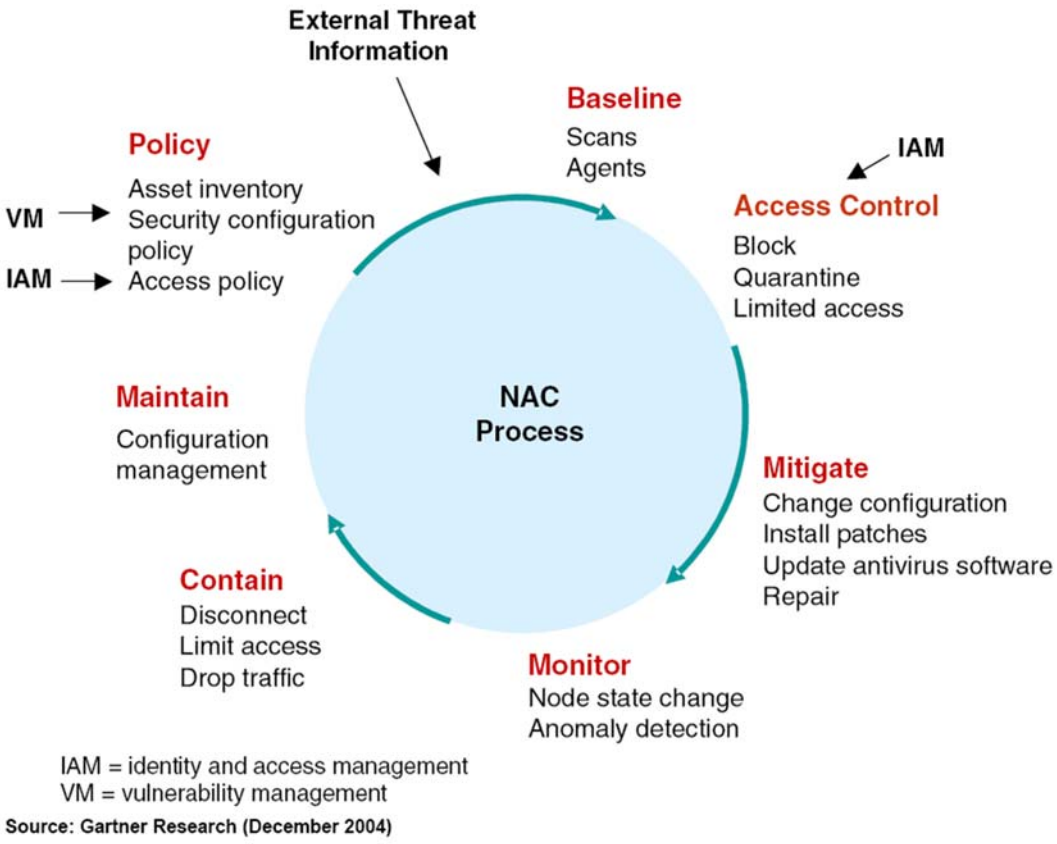

Figure 5: NAC Implementation and Management Process
(Nicolett, Orans, and Pescatore, 2004)

Policy Development

By the time you get to the NAC process, you should have a working security program in place. A security program consists of policies and supporting processes, standards, and guidelines. If this is complete, including an access policy, you can move to *Creation of Baselines.* If not, here are some tips:

1. Inventory all information assets.
2. Classify information assets according to the criticality of availability to business operations, regulatory constraints (i.e. HIPAA, SOX), and business impact if the data is compromised.
3. Create policies specifying the organization's position on secure configuration of servers, workstations, etc. At the very least, this activity must produce an access policy.
4. Create standards and guidelines to support the policies. Employees must conform to standards. Guidelines are used to provide a loose security framework upon which to build the enterprise network.

These policies, standards, and guidelines serve as input into the next process where baseline security configurations are implemented.

Creation of Baselines

Using your policies, standards, and guidelines, the next step is to define baselines. Baselines are the minimum level of security allowable when determining whether to grant an endpoint device access to your network. This is accomplished by scanning or through the use of persistent agents installed on all endpoint devices.

Scanning is typically accomplished by dynamically installing a host checker on the endpoint system. It's a good way to check devices that aren't owned and managed by your organization. The host scanner gathers information relevant to the connection policies you've implemented, and communicates that information back to the policy server. We'll step through this process in more detail in the next section.

For systems you own, persistent agents might be the right solution. These agents reside on the endpoint devices and are configured through a central management console. They keep track of the health of the system, and report status to the policy server for an access decision. Agents can provide more flexibility when your solution includes self-healing processes. For example, if the agent reports that a device is missing a patch, it might redirect that device to a restricted subnet. Once you connect to the subnet, the agent can request installation of the patch. It then notifies the policy server, which can provide access to the production network.

Management of baselines is an iterative process that incrementally improves information asset assurance as the threats against your network evolve. Clearly defined access control baselines provide input into the design and implementation of the access control infrastructure.

Access Control Design and Implementation

The design of your access control infrastructure is determined by the solution you select and the approach you take to restrict or grant access. We'll use the network design in Figure 6 as we step through three popular access methods: SSL VPN, wireless, and conference room physical connections. We're assuming a Scan/Quarantine NAC configuration. Once you complete this section, you should have a fundamental understanding of what's required to implement a NAC solution in your environment.

SSL VPN

SSL VPN is a powerful alternative to IPSec VPN. Devices designed to support this access method include host checking software to ensure baseline compliance. SSL VPN can also restrict devices to network segments, both for quarantine or general security purposes, through the use of business rules (baselines) defined in the SSL VPN device itself.

The following is an example of how an endpoint device might use NAC to gain entry to your network. This process, as well as those for wireless and conference room access, is a modified version of those found in Microsoft's *Introduction to Network Access Protection* (Microsoft, 2006).

1. The VPN client initiates a connection to the SSL VPN Termination Point
2. The SSL VPN Termination point passes the client's authentication credentials to the RADIUS server
3. The RADIUS server attempts to authenticate the user using Directory Services

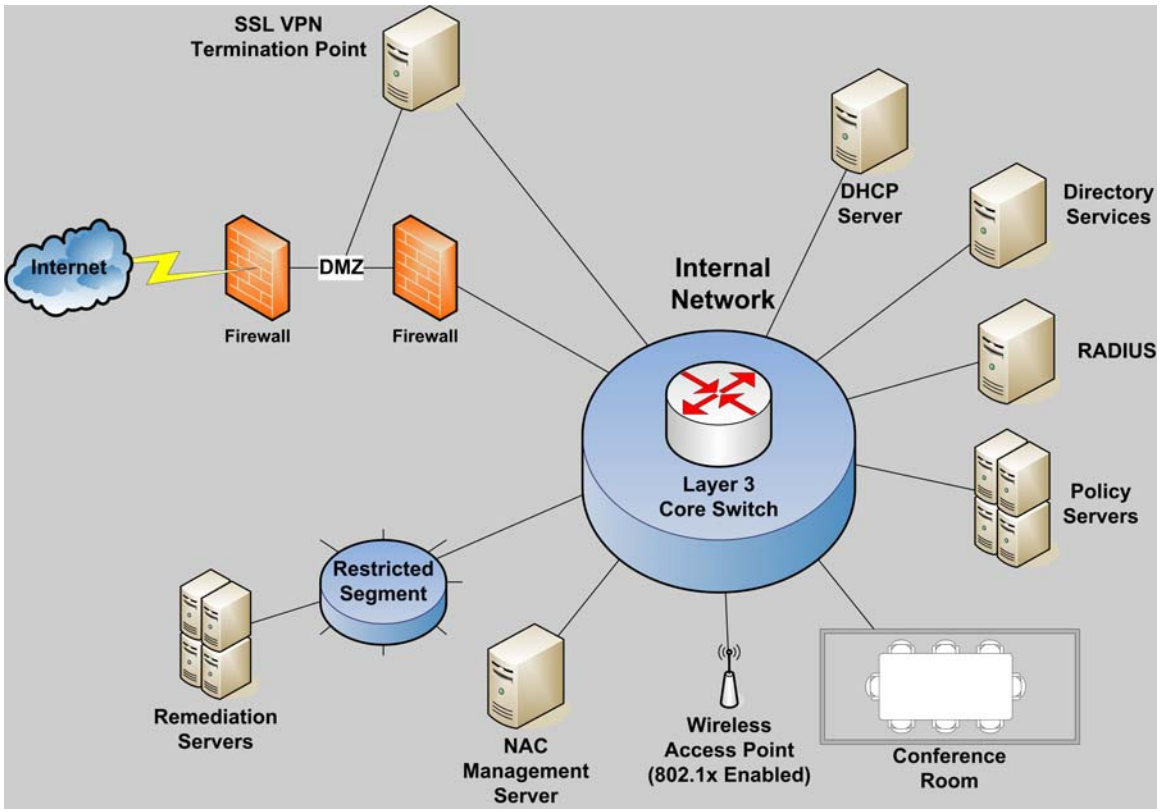

Figure 6: General NAC Infrastructure Design

4. If the authentication is successful, the NAC Management Server causes a host checker to be run on the remote device. This host checker can be an agent that permanently resides on the endpoint device, or it can be a host checking applet that is dynamically loaded when the connection attempt is made. If a host checking application cannot be run, the endpoint device is wither denied access or limited to communicating over the restricted segment.
5. The results of the host checker scan are compared to the policies configured in the policy servers.
    a. If the endpoint device meets all baseline criteria required for network access, access is granted according to applicable policies.
    b. If one or more criteria are not met, the connection is completed with access limited to the restricted segment.
        i. If configured to do so, the host checking agent on the endpoint device sends a request to the remediation servers for compliance updates.
        ii. The remediation servers provision the endpoint device with the patches, applications, configuration, etc. necessary to meet baseline access requirements.
        iii. The agent sends its updated status to a policy server. Access is then granted to the endpoint device in accordance with applicable policies.

Wireless Access

It's well known that out-of-the-box wireless connections are not the most secure. It's also well known among IT professionals that strong encryption techniques, like WPA, are required to protect data that travels over the wireless radio waves. But what about preventing rogue devices from connecting to your network through an access point (AP)? This is especially a problem with an AP that has a coverage area that, due to business requirements, extends beyond the walls of your building. And what about those high risk devices carried into your facility by authorized personnel? An 802.1x NAC solution can help meet these challenges.

An IEEE 802.1x enabled AP acts like a closed door to any device that doesn't pass authentication. This prevents potential intruders from gaining any level of access to your network infrastructure. When coupled with NAC technology, it also prevents any system that doesn't conform to the required access baselines from gaining wireless access. Let's step through the authentication and access control process.

1. The wireless client initiates a connection to the AP.
2. The AP point passes the client's authentication credentials to the RADIUS server
3. The RADIUS server attempts to authenticate the user using Directory Services
4. If the authentication is successful, NAC Management Server causes a host checker to be run on the remote device. This host checker can be an agent that permanently resides on the endpoint device, or it can be a host checking applet that is dynamically loaded when the connection attempt is made. If a host checking application cannot be run, the endpoint device is either denied access or limited to communicating over the restricted segment.
5. The results of the host checker scan are compared to the policies configured in the policy servers.
    a. If the endpoint device meets all baseline criteria required for network access, access is granted according to applicable policies.
    b. If one or more criteria are not met, the connection is completed with access limited to the restricted segment.
        i. If configured to do so, the host checking agent on the endpoint device sends a request to the remediation servers for compliance updates.
        ii. The remediation servers provision the endpoint device with the patches, applications, configuration, etc. necessary to meet baseline access requirements.
        iii. The agent sends its updated status to a policy server. Access is then granted to the endpoint device in accordance with applicable policies.

Conference Rooms

Conference rooms have the potential of being one of the weakest points of entry into your network. Unprotected physical network jacks provide a means of logical access to anyone with physical access. This is especially troublesome when conference areas are open to the public. Four common methods of securing your conference room

network access, while still allowing mobile user and vendor access, include VLAN, wireless, 802.1x, and DHCP restrictions.

VLAN - You might consider placing all your conference rooms on a single, restricted VLAN. In this way, you can provide access to resources like email and the Internet without putting the rest of your production network at risk.

Wireless - We've already discussed the use of 802.1x access controls for wireless. Disabling physical network jacks and implementing wireless access points is a good way to block unwanted network access by rogue access points, network scanners, etc.

802.1x – 802.1x isn't exclusively a wireless protocol. It's a general access control protocol that's supported by a wide variety of switches. Connecting your conference room jacks to 802.1x enabled switch ports provides the same access protection as a wireless implementation.

DHCP – Applying NAC restrictions using DHCP is similar to the process described for 802.1x and SSL VPN. With a DHCP solution, access is granted to various network segments through the assignment of IP addresses selected according to rules configured on the policy servers.

Once installed, the NAC infrastructure requires constant care and feeding. No NAC solution is perfect. Malicious packets that make it through your network portals must be dealt with by other means. This is the focus of the final four NACIMP phases.

Vulnerability Mitigation

The first goal of Vulnerability Mitigation is to ensure systems are properly patched. This also applies to NAC infrastructure components. Without application of critical security patches, your NAC solution eventually becomes window dressing for your auditors.

Be sure to review security journals and develop a relationship with your vendor technical teams. This helps keep your team up to date on current modes of attack. With this knowledge, system configurations can be modified to strengthen your defense.

And finally, ensure all anti-malware software is current. Malware that makes it past your access control safeguards must be stopped by your servers and endpoint devices.

Network Monitoring

Again, no defenses are perfect. Monitor your network for both traffic and packet anomalies. Monitoring tools should be configured to alert appropriate personnel when unusual activity is detected. This accomplishes two things. First, attack attempts in progress can be blocked before they establish a foothold in your network. Second, traffic associated with infected devices can be identified, systems cleaned, and the impact on your business mitigated.

Containment

Containment is the process of preventing high risk endpoint devices from connected to your network. Further, connected devices found to be infected or high risk by monitoring tools should be automatically moved to restricted segments. Continuous host checking is a good way to ensure systems allowed to connect remain safe.

Maintenance

NAC maintenance is closely related to Vulnerability Management. NAC policy reviews and technical team walkthroughs of the current state of the access control environment should be an active part of your NAC management process. Weaknesses or inefficiencies identified must be included in an action plan designed to maintain the effectiveness of your NAC infrastructure.

Conclusion

Network Access Control is an integral part of any layered security solution. Even if your organization isn't able or willing to spend the resources necessary for a full NAC implementation, consider an incremental approach. For example, creating restricted network segments with VLAN's is a good start. This allows you to restrict mobile and remote endpoint devices to a subset of your network, while isolating your most critical systems. Regardless of the approach you take, just staying focused on access control principles during solution design and implementation activities will go a long way toward protecting against unwanted endpoint device behavior.

 Tom Olzak, MBA, CISSP, MCSE, is President and CEO of Erudio Security, LLC. He can be reached at tom.olzak@erudiosecurity.com. Additional security management resources are available at http://adventuresinsecurity.com.

Visit Tom's **blog** at http://blogs.ittoolbox.com/security/adventures/

Listen to Tom's **podcasts** at http://adventuresinsecurity.com/podcasts

Free security training available at http://adventuresinsecurity.com/SCourses

Works Cited

Dell (2004, February). *What is VLAN routing?* Retrieved April 24, 2006 from http://www.dell.com/downloads/global/products/pwcnt/en/app_note_38.pdf

Microsoft (2006, February). *Introduction to network access protection.* Retrieved April 22, 2006 from http://download.microsoft.com/download/8/d/9/8d9b3e54-6db7-4955-9e36-58a3f0534933/NAPIntro.doc

Nicolett, M., Orans, L., and Pescatore, J. (2004, December). *Implement a network access control architecture (Gartner Research G00124425).* Retrieved April 22, 2006 from http://www.gartner.com

Nicolett, M., Pescatore, J., and Girard, J. (2004, January). *Scan, block and quarantine to survive worm attacks (Gartner Research T-21-7550).* Retrieved April 22, 2006 from http://www.gartner.com

Phifer, L. (2002). *Using virtual LANs to get more from your firewall.* Retrieved April 13, 2006 from http://www.corecom.com/external/livesecurity/vlans.htm