Controlling network access and endpoints
By David Strom

As more enterprise computing users become mobile, the chances that one of these laptops will become infected when off your enterprise network becomes more likely. And while many corporate IT departments attempt to secure their laptops with anti-virus and personal firewall software, these defenses aren't enough to keep up with the malicious software attacks that course through the Internet on an hourly basis.

So what can an IT manager do to protect their endpoint PCs? This white paper from the Trusted Computing Group (TCG) will review what options exist, show you what endpoint security does and doesn't do, and how it fits into your existing network security solutions.

1. What is and what isn't network access control

The business need for network access control is clear: Today's desktop isn't what it used to be; in fact, it may not be a desktop PC at all. As more users telecommute or connect from remote offices, it is more likely to be a laptop that is the machine of choice. These laptops are a threat to corporate networks, and change the network perimeter to one that is very porous and can be easily penetrated when the laptops are sitting in a remote coffee shop or on a hotel network. And the first thing that most of these laptops do every morning is to connect to the Internet and grab their corporate email. As the Internet becomes the main business communications pathway, companies need stronger defenses on each user's computer to ensure that they are healthy and not a potential threat to corporate computing resources.

Added to these woes is the issue that corporations need to satisfy a growing rulebook of compliance regulations that penalize corporations for personal data that inadvertently becomes public.

Network access control is a very different kind of security from beefing up your perimeter defenses with a new kind of intrusion prevention appliance, although some intrusion vendors have including various endpoint health assessment tools as part of their product offerings. It also isn't about inventing a better virtual private network, although some VPN vendors have also included endpoint assessment routines in their products. It really is an entirely new way of looking at your network and how machines connect to it, with the understanding that it isn't the user but the machine that needs to establish some form of trusted relationship.

In the past year, Microsoft, Cisco, and TCG have developed three different broad approaches for network access control. Each covers five basic issues that taken together will protect any endpoint and allow for the most robust security measures possible. These issues include:

- **Policy definition**. You should be able to set and maintain a variety of security policies for different user populations, locations and machine populations, and be able to easily modify them from a central management console.
- **Detection**. No matter how your users connect to your enterprise network, your system should be able to detect them. This includes using agents or agent-less operations on each client; no matter what operating system version they are running, and whether they are on the local headquarters' network or accessing it remotely.
- **Health assessment**. Your security system should be able to scan the endpoint and determine compliance with your policies. Ideally, the scans should take place prior to any network access, but your system should also allow other checks to occur after login too, such as which Web browser the user is running and what other security policies should be applied.
- **Enforcement**. Your policies determine what network resources should be protected, including switches, VPNs, servers, and so forth. You should be able to quarantine resources or refuse network access entirely, depending on what policies you maintain.
- **Remediation.** If clients don't pass muster, what happens? The ideal system should kick off anti-virus signature updates, or apply patches to the OS, or other measures after they have been quarantined, so that users can eventually connect to the corporate network after everything is brought up to date.

Obviously, this is a tall order and most of the network access control approaches don't offer complete solutions for all five issues, or for many different types of operating system and browser combinations. And there is a lot of confusion surrounding network access control (see Sidebar 1 on the four myths of endpoint security).

Let's see how these solutions fit in with what you already have.

2. How does your existing security infrastructure work in terms of protecting endpoints?

Before you can implement any network access control solution, you first need to understand what you already have in terms of security infrastructure. Part of the problem is that most of this gear has been designed to authenticate users, not devices, to your network. That is why we are in the situation today with virus

infections and rootkits, because they are launched by PCs that have already been authenticated and are inside the corporate network, even though they are anything but benign.

Some of the endpoint access control products come with their own intrusion detection/prevention systems, virtual private network gateways, RADIUS servers and firewalls, and these elements could duplicate what you have in place already. The bonus is that they have extended these traditional security products to include some form of endpoint scanning and enforcement.

If you don't have these elements, it may sound appealing that the network access control system rolls these protective measures into their solution. But this is worth further investigation, since the fine print may reveal that you still have something that is incomplete. For example, if your corporation doesn't yet have any VPN, buying one that is included in an endpoint solution may mean that the local network PCs aren't scanned for endpoint health issues, which brings us back to square one in dealing with this problem.

And it is also important to understand that just because a particular vendor supports someone else's VPN or IPS doesn't mean that this support is seamless or that there is a single repository of security policies that can be applied across local and remote users. Before purchasing any solution, check to see how security policies are stored and whether they will cover other situations outside of any endpoint health assessment.

3. What architectural decisions does an IT manager need to make in terms of network access control deployment?

Part of the critical deployment of any endpoint solution is in understanding where in the network the various pieces of that solution will be placed, and whether your existing network architecture needs to be modified. There are two basic schemes that the network access control solutions support**: in-line and out of band**.

When network access control solutions operate in-line, any network resources located behind them will be protected and only healthy network clients can pass through and gain access to these resources. This is similar to how a VPN operates for remote access, and has appeal for those networks that have placed or who can place all their critical servers and network resources on a single subnet.

When the network access control solution operates out of band, it is typically connected via a span port, watching over all of the network traffic on that particular subnet. Then they insert themselves into the network stream once a user has successfully authenticated through an Active Directory or VPN login. Some vendors support both attachment methods.

Choosing the right architecture is also a matter of relative throughput that each device will be able to handle passing through it. Some products are designed to operate at gigabit speeds and be placed at the network core, while others are more appropriate closer to the network edge.

4. What about desktop deployment decisions?

Each network access control product uses a combination of software agents to scan the endpoint and make determinations about its health. This software also may act on the health information for remediation of the endpoint and protecting network resources.

There is a lot of work for these agents to do. They have to examine the open ports and running services, look through the file system for any rogue files or malware, monitor Windows registry keys, keep up to date on anti-virus signatures, and check whether a personal firewall is running or has been tampered with. With some products, all this scanning and checking can extend the login process by several minutes, which could flummox users who were expecting something quicker and might reach for the phone to call tech support before the agents let them login.

There are three basic types of agents that potentially can be used by each appliance:

- A **"thick" agent** that is permanently installed executable file on each endpoint PC,
- An **on-demand agent** that doesn't persist beyond the period of time that a PC is connected to your network, typically delivered by a browser session or part of the network login process, and
- An **agentless solution** that doesn't place any software on the endpoint, but operates with something that already exists on the PC.

Let's go into more detail on all three types of approaches, because the subtleties are significant and can determine how you'll end up using each product. There are differences in terms of what kind of software is delivered to each endpoint, how long it sticks around, when it does the health assessment during the login/connection process, and whether it can "dissolve" or remove itself automatically after the endpoint is no longer connected to your network. In many respects, the thick and on-demand agents are conceptually similar to the difference in VPN approaches between IPsec and SSL clients.

The thick agent is the easiest to understand, but the hardest to universally deploy, because it means that IT must be able to manage and control these systems. In most cases, the thick agent requires that each desktop allow

administrator access for Windows, because you are installing software on every machine. Most of the vendors offer a thick agent for Windows 2000/XP and some either currently or have announced support for Mac OS. Obviously, depending on the mix of clients will determine your strategy here.

The thick agent is usually able to do the widest assessment of endpoint health, including scans of the Windows registry, file system, system services and open ports to determine if the machine has been compromised. It is also able to perform the remediation tasks to bring the endpoint back into compliance.

The trouble comes for unmanaged PCs that fall outside the control of IT, such as those being used by guest workers or business partners. This is the hairy edge of network access control, and the one that you'll need to scrutinize carefully to understand the differences among the various vendors. Some of the thick agents require administrative Windows access for their installation.

The issue is an infected machine may be able to do damage on your network prior to any login, just by being on a particular network segment. Most of the vendors have this situation covered by doing some kind of rudimentary health assessment prior to providing an IP address or other network credentials to a endpoint.

On-demand agents can cover the unmanaged PC because the agent doesn't get delivered to the endpoint until the machine tries to connect to your network. These typically take the form of a JavaScript or Active X control (or both) that are loaded via a Web browsing session at the time of connection, similar to how SSL VPNs work with their browser-based clients.

Some of the on-demand agents actually clean up after themselves, what the industry is now calling "dissolvable" agents. Some vendors take things a step further by placing each endpoint on its own private VLAN, thereby ensuring that risky devices remain isolated from others.

These on-demand agents are specific to particular operating systems and browser versions, and while most support Firefox and IE browser, there are some exceptions, particular as you move away from Windows OSs. Other vendors have been slow to support non-Windows systems too, so if you have a significant population of these machines, examine when they actually added support for the appropriate OS.

Finally, there are the agentless approaches, which sounds peculiar but is actually quite clever, because they leverage something that is already present on the endpoint that the security system can query. Typically, these don't offer much in the way of remediation, but do offer some minimal identity protection. Again, these are OS specific pieces of software, and in some cases tied to specific OS version levels too.

Some vendors offer multiple agents but with different capabilities. For example, some vendors offer an Active X control for Windows that will perform its health assessments but have different features for non-Windows users that can only do minimal identity controls. Others offer on-demand agents on Mac OS and Linux, and have thick agents for Windows 2000 and XP.

5. What is the TCG doing for network access control?

So far in this paper we have kept our discussion about network access control fairly general and not really gotten into specifics about particular products. There are three major architectural approaches and vendors have aligned themselves with one or more of these camps with their product offerings. Sidebar 2 summarizes the efforts underway by both Microsoft and Cisco with their architectures, and in this section we'll talk about what we at the TCG are doing.

Trusted Computing Group's Trusted Network Connect (TNC) is the only open standards-based network access control effort, and it has standards for just about every aspect of endpoint security either published or in early form. Unlike the Microsoft and Cisco architectures, TNC is not tied to any particular vendor's hardware or operating system.

TNC allows IT administrators to define one or more security policies and check each network-connected system for compliance with these policies when the endpoint connects to the network. It is designed to work with a wide variety of network equipment so it does not require massive equipment upgrades or fork-lift replacements. And it is focused on non-proprietary vendor-agnostic technologies that are tested for interoperability among many vendors to ensure a multiple-vendor approach to network access control solutions.

There are two important parts of TNC that while both optional can be used to build very sophisticated endpoint security systems. The first is the Trusted Platform Module (TPM), a standards-based encryption chip that is present in a growing list of new PCs. These chips can serve as the basis for a known trusted state of the PC, and be exploited by a variety of software tools to ensure that the endpoint hasn't been infected with a virus or rootkit. The TPM securely stores digital keys, certificates and passwords. Because it is security hardware, it is especially difficult to attack virtually or physically. TPMs are sold complete with internal firmware so that the chip does not have to be programmed. The TPM can be used for a variety of purposes, including trusted software downloads, secure network communications, reliable identification of system peripherals, and secure local storage. See this white paper for further details: http://www.trustedcomputinggroup.org/groups/tpm/embedded_bkgdr_final_sept_14_2005.pdf

For a list of which PCs contain the TPM, see Tony MacFadden's chart here: http://www.tonymcfadden.net/tpmvendors.htm

6. What is the relationship of 802.1X to TNC?

The second optional component of TNC is support for the IEEE 802.1X authentication protocols (802.1X for short). When used together, 802.1X and TNC provide strong network access control for wireless and wired Ethernet networks. Of course, TNC can be used with other network access methods but 802.1X is an especially strong and popular one.

802.1X is the standard way to control access to a wireless or wired Ethernet network. Network access control is especially important in a wireless network, which is potentially open to anyone within radio range. But wired networks can also benefit from the strong protection provided by 802.1X.

When a machine connects to a network protected by 802.1X, it is immediately prompted by the wireless access point or wired switch to authenticate itself. The authentication exchange is forwarded to an authentication server, which decides what sort of network access should be granted (if any) and informs the wireless access point or switch.

Two aspects of this authentication exchange are especially noteworthy. First, no network access is granted until the exchange is complete. Thus, the protection provided by 802.1X is especially strong. Second, the authentication exchange uses the Extensible Authentication Protocol (EAP) so the exchange is almost infinitely flexible and extensible. The Trusted Network Connect standards employ this extensibility to verify a machine's health as part of the 802.1X exchange.

Wireless access points and switches that support 802.1X don't need to be modified to support TNC health checks. They simply pass the health information on to the authentication server, such as a RADIUS or Active Directory server, which is extended to support this new dialog. Because the health information is implemented using EAP, it just looks like another authentication method. The authentication servers need to be upgraded to understand this additional information.

The 802.1X authentication protocols aren't a necessary part of TNC, but if they are used they can support a complete range of endpoint health assessment tasks. What is new about this method is that 802.1X is being used across both wired and wireless connections for both local and remote users to identify not just whether a user has the appropriate network login credentials, but also whether the user's PC has the appropriate endpoint health credentials to participate on the network.

7. What are the specific TNC standards?

The charts below summarize the various standards efforts that are underway as part of the TNC effort and their current version. We also provide some of the typical vendors who have announced products that meet the particular standard.

TNC Standards Summary

Standard	Version	Purpose
TNC Architecture	1.1	Overall architectural summary
TNC IF-IMC: Integrity Measurement Collectors	1.1	Client plug-ins that report on endpoint health.
TNC IF-IMV: Integrity Measurement Verifiers	1.1	Server plug-ins, compare client state reports with policies.
TNC IF-PEP: Policy Enforcement Points, Protocol Bindings for RADIUS	1.0	Specifies the integration with PEPs such as VPN concentrators, switches and wireless APs – and how messages are sent between AAA servers and PEP to allow/deny or grant limited access.
TNC IF-T: Protocol Bindings for Tunneled EAP Methods	1.0	The various transports using EAP such as TTLS, FAST and PEAP.
TNC IF-TNCCS	1.0	This specifies interoperability between the TNC client and server and how they exchange messages and can be made transport protocol independent.
TNC IF-PTS: Platform Trust Services	1.1	Defines how the TPM can be used with overall TNC architecture.

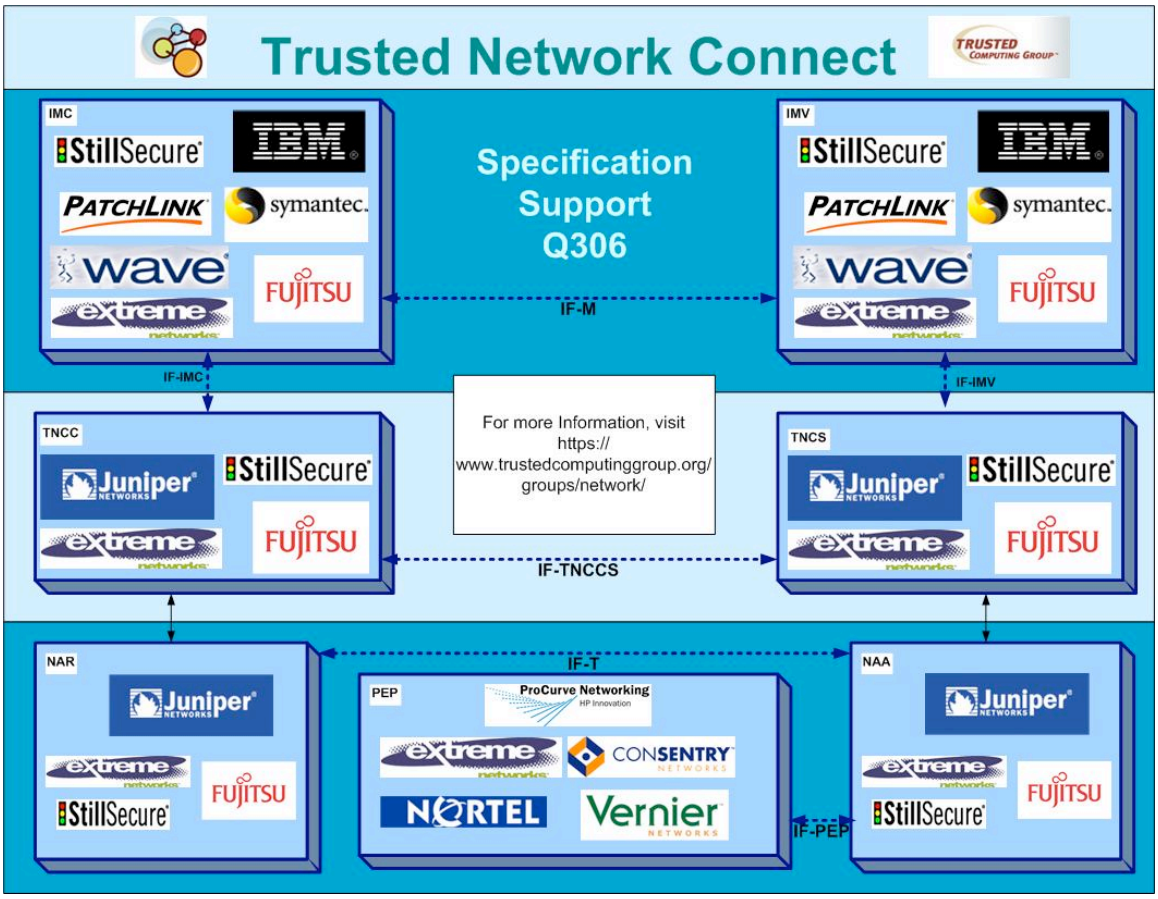

Figure 1. TNC architecture, standards, and typical products

8. Conclusion

Network access control and endpoint security is a rapidly emerging field with lots of players, evolving standards, and new products announced weekly, or so it seems. TNC is the only open-standards based effort that is underway to solve this problem, and this paper presents an outline of what are the issues in implementing secure endpoints and what a typical corporate IT manager will need to understand before choosing any particular vendor or more likely, multiple vendors.

Sidebar 1: Four myths about endpoint security

There is a lot of misinformation about network access control and endpoint security. Here are the four biggest ones that need some further explanation.

Myth #1. Every endpoint can be secured with the same software solution.

Not true. If you have the unfortunate circumstance to have a heterogeneous network composed of many different operating systems, embedded devices with their own IP addresses, and switches from three different vendors, you are in for some trouble. The problem is that the software agents that root out the dirty work are very specific in terms of browser version and operating system. Some require initial administrator rights to be installed on the endpoint, while others are only present during the time an endpoint is logged into the network and disappear or "dissolve" when the session is terminated.

Then you have the situation of embedded devices. How do you protect those? Some network access control solutions have "white lists" where you can specify that your Web cameras and print servers don't have to be scanned for them to attach to your network. That is all well and good, until the bad guys figure out a way to compromise these devices and then trouble begins. Some vendors have ways to monitor these embedded devices as part of their systems, so if you have a lot of them on your network you might want to start with them first.

Before you examine a product, understand the device protection portfolio that is offered by each vendor and also what they have promised for the coming year. Many vendors are still lacking basic Mac OS support, for example.

Myth #2: It is easy to identify an unhealthy endpoint and block it from coming on your network.

Again, troublesome. Figuring out what ails that endpoint isn't simple. You need some sort of scanning routine – and done regularly – to determine if its file system, registry keys, anti-virus signatures, OS patch levels, and whatnot are up to spec.

As you might imagine, all this scanning takes a bunch of time during the login process, so setting up end user expectations is critical before those support calls come in that "I can't login to the network."

Some of the products require all kinds of specific information to check, such as particular anti-virus signatures or personal firewalls. Some products scan at different degrees of depth depending on what kind of agent is doing the scanning. It pays to check this out carefully before proceeding further.

Myth #3: It is easy to cure an unhealthy endpoint.

Also false. Any remediation needs to cover all the things that can go wrong, and do it as automatically as possible. Otherwise, your end user support lines will light up like a Christmas tree the moment you turn on any solution. Some of the products are better at remediation than others. Some, like Cisco's solution, just send an unhealthy endpoint off to the network equivalent of Siberia without trying to fix the problem.

Myth #4: Your network access control solution will work seamlessly with your VPN.

It makes some sense and certainly would be nice if that creaky old VPN that you have had around for several years could just interoperate with that nice new shiny network access control system. But no, this is not to be. For those of you that have enterprise IPsec VPNs, you are in a better place to implement an network access control solution, provided that you run those secure IPsec protocols on all of your local machines too. Most of the endpoint products support this approach, however cumbersome and unattractive it sounds at first.

If you aren't willing to completely re-architect your LAN logins, and don't yet have a VPN that you want to keep using, some of the vendors offer network access control solutions that integrate VPNs into their package.

Sidebar 2: Microsoft and Cisco approaches to endpoint security

There are three major architectural approaches that are currently being developed to specify endpoint security, with a fourth in the works to try to bring some common ground. The three approaches differ in terms of what they offer to protect, how they implement assessment, remediation, and enforcement, and have a different set of vendors supporting each approach. While there are other resources [LINK] that go into more details, let's quickly summarize h and show the major differences between what Cisco and Microsoft are trying to do.

Cisco's Network Admission Control (NAC) takes the approach to control the network access layer through trusted modules that are implement in the router and switch code, working together with agents that are running on Windows and Linux endpoints. Some analysts feel that NAC will require too many pieces to deploy properly. You'll also need to ensure that the IOS firmware in these routers and switches are brought up to date, since NAC only supports the most current versions.

Cisco's architecture is short on remediation – there is little to manage patch levels on the endpoints, and it has ceded this territory to Microsoft in its most recent announcements. There is also little in the way of granularity for different levels of health – either an endpoint passes muster or it is shunted off to some VLAN with limited access until it can be fixed.

In this paper, we use the term "network access control" to indicate the general approach that many vendors are taking in terms of improving endpoint security. We use the acronym "NAC" to designate Cisco's own approach.

Cisco has a second approach to network access control that it acquired from Perfigo and now calls Cisco Clean Client Access. It is offered as both a software and hardware appliance, and works with Windows-based agents.

Microsoft's Network Access Protection (NAP) has yet to be implemented in any product, although agent pieces of it can be found in Windows Vista, and antecedents in the VPN quarantine features found in Windows Server 2003. Microsoft will support XP and include NAP in Longhorn Server, which is still in beta. It brings security policy and enforcement to Windows. Unlike NAC, it has multiple remediation paths that are part of the architecture.

Microsoft is concentrating on the more recent versions of Windows in its support of NAP, and most certainly older operating systems will need to run multiple agents to participate in this architecture. They and Cisco have announced various interoperability statements that combine products from both companies.

There is actually a fourth effort underway, by the IETF, to identify common protocols that are used in all three architectures and use its consensus process to produce interoperable interfaces that are common and can be used to communicate amongst the various NAP, NAC, and TNC products.

Sidebar 3: How the So. Carolina government is using TNC

One example of a TNC implementation is in Columbia, S.C. at the state Department of Probation, Parole, and Pardon Services, the government entity that manages many aspects of offenders re-integration into society once they are released from the Department of Corrections. The department has more than 50 remote offices around the state and supports agents in more than 40 different local courthouses that need to be connected to its network, along with hundreds of field agents that are traveling around the state working with clients.

The department recognized a need for network access control, given that many of its endpoints are at risk based on connecting from various secure and insecure computing environments. They operate a mixed network infrastructure of Juniper routers and firewalls, various security appliances, and HP switches, with a few remaining Cisco routers too. All of this supports a wide variety of desktops, tablets, and servers running several different Windows, Novell, Linux and Unix operating system versions. They use McAfee's Enterprise Total Protection Suite including Entercept and McAfee's Policy Enforcer for their desktop access control, protection, host- intrusion prevention, and anti-virus scanning.

They got involved with the TCG after implementing a whole disk encryption solution last year, using the SafeGuard product that can store the encryption keys on the TPM chipset. "Standards are the only way to go," says David O'Berry, the IT director at the department. "Everybody wants to

be Cisco or Microsoft whether they want to admit it or not. But there is no way that you are going to maintain a homogeneous environment, and I can't be hamstrung by being an all-Microsoft or all-Cisco environment. For that matter, I cannot be held hostage by any single vendor because complexity is only going to go up as we continue to implement our network access control solution making TCG compliance all the more important at every juncture."

O'Berry is a big proponent of using standards-based products to ensure the widest choice of the best products. "If you get the standards right, then the vendors can compete on the basis of product features. I don't want to have to buy a mediocre product from a vendor now because they have me over a barrel," he says.

The department hopes to have its endpoint security solution in place by mid-year and has worked and continues to attempt to work only with TCG compliant vendors in trial or beta modes. "You have to truly protect the people that use your network or the business will suffer a staggering cost going forward," he says. "If you can get these endpoints squared away and get some kind of stability and security to the edge of the network, you have a chance at winning the battle. We think TCG is on the right path and hope that more vendors include support for the TNC standards in their products."

For more information, please contact:

Trusted Computing Group Administration
5440 SW Westgate Drive Suite 217
Portland, OR 97221
Email: admin@trustedcomputinggroup.org
Phone: (503) 291-2562, Fax: (503) 297-1090