



Controlling Network Access and Endpoints

November 2007

Trusted Computing Group Controlling Network Access and Endpoints

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.

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. 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 firewall, although some firewall vendors have included 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 just the user but also the machine that needs to be trusted.

In the past few years, a variety of vendors have developed different 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.

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 whose users have already been authenticated and granted access to the corporate network, even though their PCs 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.

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

Microsoft's approach to Network Access Control

Microsoft's network access control architecture is called Network Access Protection (NAP). Windows Vista includes a NAP client. Additional NAP support will be included in Windows Server 2008 and Windows XP SP 3.

In order to ensure maximum interoperability and compatibility for the NAP architecture, Microsoft has endorsed the TNC architecture and critical TNC standards. In particular, the client-server protocol used by NAP is a TNC standard: IF-TNCCS-SOH. This allows customers to mix NAP clients, TNC clients, NAP servers, and TNC servers interchangeably, as long as they support the IF-TNCCS-SOH standard. This removes any question about Microsoft's commitment to the TNC standards and to interoperability in this critical area.

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.

Out of band network access control solutions employ a variety of different techniques. Most often, they integrate with existing in-band components (routers, switches, etc.) so a control node receives notice of new endpoints. The control node scans and/or monitors the endpoints. Non-compliant endpoints are isolated using the in-band components and remediated. Some vendors support both in-line and out of band attachment methods.

Choosing the right architecture is also a matter of determining the 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.

What about desktop deployment decisions?

Many network access control products use some form 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 are three basic types of agents that potentially can be used by each appliance:

- A **"thick" agent** that is a 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 as 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 or scans the endpoint remotely.

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, what sorts of health assessments it can do, 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 XP and Vista. Some either currently support or have announced support for Mac OS. Obviously, your 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, interactions with anti-virus or other client security software, and

Cisco's approach to Network Access Control

Cisco has taken a more proprietary approach to network access control. Both of their products in this area (the Network Admission Control Framework and the Network Admission Control Appliance) are built on proprietary protocols and APIs. Therefore, they only work with Cisco equipment and specifically licensed partners.

Fortunately, Cisco has realized that this proprietary approach won't work in the long run. They are not willing to join TNC effort but TCG members are working with them in the IETF to agree on standards. As with most IETF efforts, this is not a quick process but it should eventually result in a reconciliation between Cisco and the rest of the industry, which has converged on the TNC standards.

other measures to determine if the machine has been compromised and whether its defenses are up to par. 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 guests, customers, contract 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. Few guests will want to install a thick agent on their machine.

On-demand agents can cover the unmanaged PC because the agent is lightweight and doesn't get delivered to the endpoint until the machine tries to connect to your network. On-demand agents typically take the form of a Java or ActiveX control that is 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.

Finally, there are the agentless approaches, which sound peculiar but are quite clever, because they leverage something that is already present on the endpoint that the security system can query. These are limited in their capabilities to whatever can be done with the built-in software on the endpoint. Agentless approaches are the easiest to employ because no extra software installation is needed. They may eventually become the most common approach as endpoint vendors build more sophisticated network access control capabilities into their systems.

Some vendors offer multiple agents but with different capabilities. For example, some vendors offer an ActiveX control for Windows that will perform exhaustive health assessments and a Java agent for non-Windows users that can only do minimal controls.

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 or architectures. Now it's time to take a serious look at the major efforts under way in network access control.

Historically, there were many incompatible network access control architectures. Each vendor brought their own products to market and none of them could interoperate. However, this has changed in the last few years. Customers and vendors have turned away from closed network access control systems, realizing that open standards are necessary to ensure that all endpoints on the network can be health checked.

The Trusted Network Connect (TNC) is an open architecture and a set of open standards for network access control. TNC was created by the Trusted Computing Group (TCG), an industry consortium that creates open security standards. Using the TNC standards, customers are free to build their network access control system with products from any vendor. Products that support the TNC standards have been tested to ensure that they will work together. This compatibility is the reason that vendors and customers alike have flocked to the TNC standards in recent years.

TNC allows IT administrators to define one or more security policies

How a government agency in South Carolina 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. The network infrastructure is a mix of Juniper routers and firewalls, various security appliances, and HP switches, with a few remaining Cisco routers too.

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, as did some proprietary approaches.

There are two important but optional parts of TNC that can be used to build very sophisticated endpoint security systems. The first is the Trusted Platform Module (TPM), a hardware security chip or function based on TCG standards that is present in all currently shipping corporate-grade laptops and desktops. Because it is a hardware security component, the TPM can securely establish the state of the endpoint, ensuring that the endpoint hasn't been infected with a virus or rootkit. For a more detailed discussion of TPM, TNC, and rootkits, see https://www.trustedcomputinggroup.org/news/Industry_Data/Whitepaper_Rootkit_Strom_v3.pdf.

In addition to its role as an optional component of a TNC system, the TPM can be used for other purposes. For example, it can securely store private keys and passwords. The TPM can also be used for a variety of other purposes, including trusted software downloads, secure network communications, reliable identification of system peripherals, and secure storage. Because the TPM is security hardware, it is more secure than software and designed to resist many attacks. TPMs are sold complete with internal firmware so that the chip does not have to be programmed. See this white paper for further details: http://www.trustedcomputinggroup.org/groups/tpm/embedded_bkgdr_final_sept_14_2005.pdf

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.

What are the specific TNC standards?

Figure 1 (below) summarizes the various standards included in the TNC architecture. More information on these standards is available on the TNC web site: <https://www.trustedcomputinggroup.org/groups/network>

TNC Standards Summary

Standard	Version	Purpose
TNC Architecture	1.2	Overall architectural summary
TNC IF-IMC: Integrity Measurement Collectors	1.2	Client plug-ins that report on endpoint health.
TNC IF-IMV: Integrity Measurement Verifiers	1.2	Server plug-ins, compare client state reports with policies.
TNC IF-PEP: Policy Enforcement Points, Protocol Bindings for RADIUS	1.1	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	1.1	The various transports using EAP such as EAP-TTLS, EAP-FAST and PEAP.
TNC IF-TNCCS	1.1	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	1.0	Defines how the TPM can be used with overall TNC architecture.

Figure 1. The TNC standards

Which vendors and products support the TNC standards?

Figure 2 shows vendors that have implemented or announced plans to implement the TNC standards. Sidebar 1 highlights the approach that Microsoft has taken to the TNC standards.

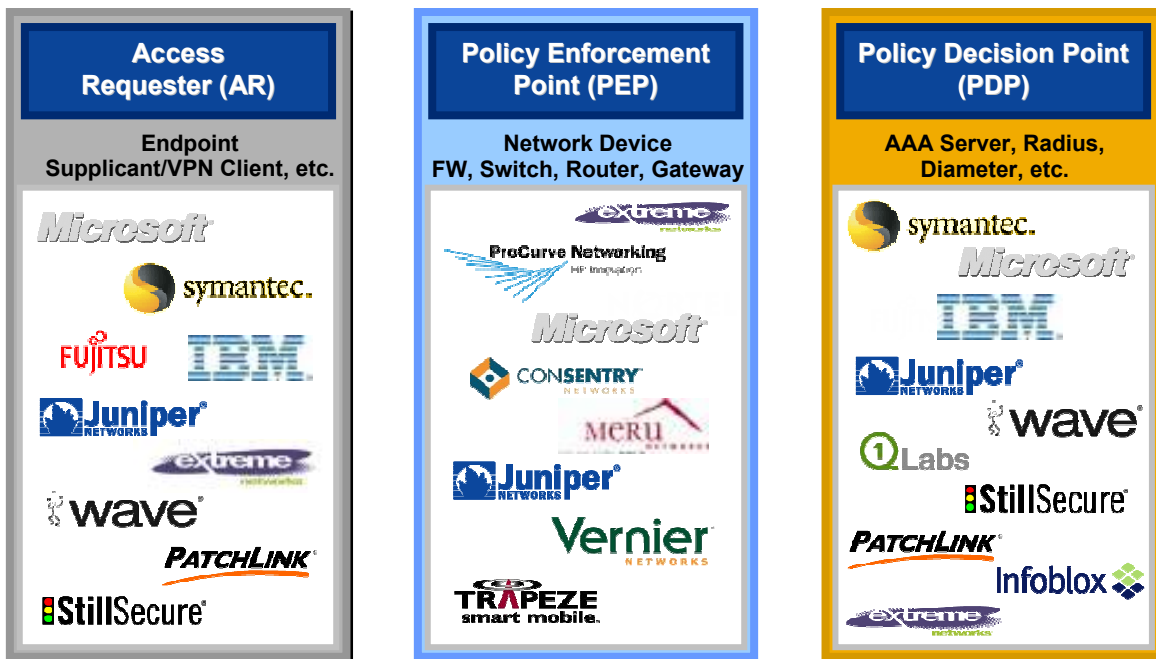


Figure 2. Vendors who support the TNC standards

Conclusion

Network access control is a rapidly emerging field. New products, alliances, and standards are announced weekly, or so it seems. This paper presents an outline of issues that IT managers should consider in implementing network access control. It also describes Trusted Network Connect (TNC), a full set of open standards for network access control.

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 Enterecept and McAfee's Policy Enforcer for their desktop access control, protection, host- intrusion prevention, and anti-virus scanning.

They became familiar with the TCG specifications 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

3855 SW 153rd Drive

Beaverton, OR 97006

Email: admin@trustedcomputinggroup.org

Phone: (503) 619-0562, Fax: (503) 644-6708

www.trustedcomputinggroup.org