Agentless Endpoint Security

Cambia

Cambia™

Stronger Security. Constant Compliance.

Abstract

Endpoint Security (EPS) has become the new standard by which the IT security industry flexibly defines "network perimeter security." The resulting flurry of companies offering agent-based endpoint security products, often with or without an accompanying centralized IDS/IPS, has fueled debate between the agent-based vs. agentless endpoint security camps. This white paper examines the agent-based EPS approach and offers insight into alternative agentless approaches and what to look for when considering "going agentless."

Table of Contents

Introduction

"Technology...is a queer thing," said C.P. Snow. "It brings you great gifts with one hand, and it stabs you in the back with the other."

Thus it is with the world of networking today. Advances in technology have made it easier for people to connect to networks via wired or wireless technology at the office, at home, at the airport, or even in a neighborhood restaurant. But the ability to connect to a network brings its share of perils to the network owner—what if the connecting device isn't up-to-date with antivirus software or the latest operating system patches? What if vulnerable ports are open? What if its firewall is turned off? What if it's already infected with the latest virus? Unauthorized devices connecting to a network can inflict significant damage in no time flat—if allowed to join.

Perimeter security has become an even hotter topic recently because the definition of the "network perimeter" has changed so dramatically. Being able to investigate, quarantine, and repair a device before it fully connects to a network is key to maintaining a strong security posture. Also, as part of some public and private regulatory requirements, new hosts must undergo compliance scanning before being allowed to connect to networks to prevent information theft.

It's no surprise then that endpoint security, or EPS, is becoming an increasingly important component of an enterprise's network security infrastructure. But the method of employing EPS on a network is still up for debate. Is an agent-based or an agentless approach better? In this whitepaper, we'll cover how both solutions operate, the benefits of the alternative agentless approach and what to look for when considering "going agentless."

Agent-based Endpoint Security

History is full of famous–and infamous–agents. There were secret agents like the Rosenbergs, a nice-looking all-American couple convicted of stealing U.S. atomic secrets and selling them to the Soviet Union. The little-known Dusko Popov, codenamed "Tricycle," made his mark by feeding the Soviets errant MI-5 data during WWII and was purportedly the basis for Ian Fleming's James Bond character. (Fleming, then working for the British intelligence, was charged with keeping tabs on Popov.) For good or for evil, agents have made names for themselves in the annals of history.

In the technology world, some of the biggest names in the business have also paved the way for themselves through the use of agents—sometimes dozens of agents. Companies such as Microsoft and Cisco have set the stage for the deployment and widespread use of agents and agent-based applications in everything from e-mail applications to data mining.

Endpoint security (EPS), recently heralded as the next coming in network defense, is no exception, as it often requires agents. The theory behind agent-based EPS is that agents installed on hosts can monitor compliance on the host and serve as keys to the gatekeeper on the network, while ensuring that assets comply with defined security policies. While the theory is sound and is backed by some of the biggest hitters in the industry, there are limitations inherent in this design.

Agent-based EPS purveyors would have the public believe that having an agent on every network asset will solve all security woes and that the past problems with agents in general have been solved. The glossy brochures look nice and it would be a nice tale to believe. Even so, forgetting for a moment the arguments against agent-based solutions (the resource usage, the difficulties in managing updates, and the inter-department corporate wrangling that ensues when the topic of agents comes up), one undeniable truth remains: an agent can only work if it exists.

So, while it's great from an inventory standpoint to know how many assets (with agents) logged onto the network yesterday, and how many of those had up-to-date antivirus software installed, it's not always the known assets that cause the most problems.

It's what you ***don't*** know about that will get you.

Agent-based Endpoint Security in Action

Before discussing relative advantages and disadvantages, it's important to understand how agent-based EPS works. Let's take a simplified look at how EPS works when a new host tries to connect to the network:

Step 1: A new asset, or host, connects to the network and requests access to network resources using DHCP and ARP requests.

Step 2: The network access device (NAD) sends a request to the access control server (ACS) for posture information.

Step 3: The ACS requests credentials from the agent on the new host.

Step 4: The host agent gathers information on the host and sends it back to the ACS where an access control list (ACL) is applied and the decision to allow, limit or quarantine the host is sent to the NAD for enforcement.

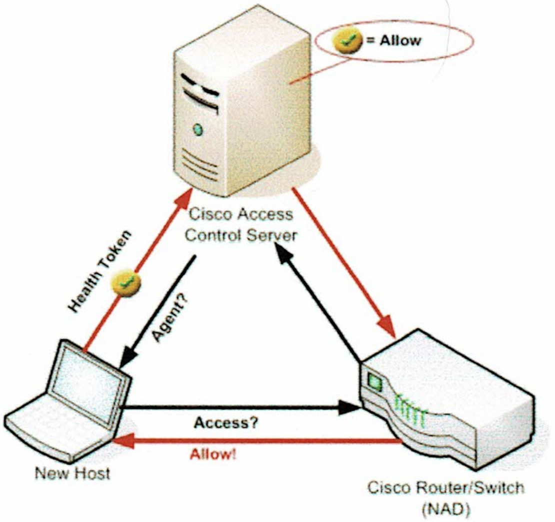

In the illustration above, the new host had an agent and was able to provide sufficient information to warrant a "Healthy" token and was allowed access to the network. As an example, in Cisco's NAC schema there are six posture token levels:

- **Healthy** (in compliance)
- **Checkup** (meets minimum compliance, but needs updates)
- **Transition** (a temporary status, access usually limited)
- **Quarantine** (out of policy, limit access to remediation sections of network)
- **Infected** (Danger! Immediately quarantine)
- **Unknown** (unable to determine compliance level, quarantine until status can be determined)

If the host is unable to provide the proper credentials, the NAD applies a quarantine access control list (ACL) to the host. Depending on the ACL, the NAD can limit access, giving the host access to certain servers, or it can completely isolate the host until it meets compliance requirements.

For the sake of brevity, let's assume that in order to meet yearly security audit mandates, most corporate security requirements are built around zero-tolerance when it comes to compliance, and in order for a host to gain access to even the most rudimentary of network services, it must first meet these full compliance requirements. When an outside host attempts access, this requirement is usually fulfilled either by requiring that the host install a security agent, or by putting the host in temporary quarantine but allowing access to a remediation server where the user can perform the necessary corrections such as applying certain patches and running an antivirus scan using the latest .dat files.

Agent-Based EPS Solutions are Limited

The Agent as a Source of Security Risk

In the previous scenario, the most common method of ensuring the new host meets compliance is by downloading the trusted agent. In theory, this is also the simplest solution, but this path to compliance brings up its own challenges:

- What do you do when the new host belongs to a third party with a strict policy against installing unapproved software on the system?
- What kinds of problems are created when the IT security group wants to put agents on systems supported by other IT groups?
- How do you prevent users from tampering with the agent's settings, or shutting it off completely?
- What if your vendor's agent doesn't support all of your platforms, such as Solaris or Linux?

The reality is that agents are inherently untrustworthy. They can be disabled, misconfigured, and compromised by user actions that range from the accidental to the malicious. Relying on a system to report reliably on its own security status is a bit like asking students to grade their own tests. Don't be surprised if they all pass.

Beyond Pre-Admission Scans – What Next?

But wait. There's more. Once a host passes the initial hurdle of the pre-admission scan and has been granted access to the network, it moves from being an external threat (a rogue asset), to being an internal asset that has achieved some minimal level of trust. The expectation here is that basic precautions have been taken:

- Antivirus is up-to-date
- Firewall is turned on
- Vulnerable ports are turned off
- Critical security patches have been installed
- Vulnerable applications such as Instant Messaging have been removed or turned off

After passing the initial compliance hurdles, agent-based EPS typically doesn't do much more. That's because agent-based EPS solutions tend not to run continuously. They are "admission" solutions, not "continuous" solutions. With agent-based EPS, much attention is focused on making sure the bad stuff doesn't get in, and that's good; but where the truly paranoid IT security expert earns her pay is by looking at the entire security picture and asking, "What next?"

What happens after the host has met initial compliance objectives, and is allowed into the network? What happens if an infected USB drive is then inserted into the host, or the contractor using the new host opens a network file share with wide-open read permissions and that file contains sensitive or private information?

These are actions that are probably benign, but can signify possible problems. And while IDS and IPS are great things to have, they don't typically do much good until after the threat has materialized and is trying to do whatever damage it is designed to do.

The bottom line is, if an agent-based EPS is not running continuously, once the host has been given access to the network, what measures are taken, if any, to ensure the host retains the level of compliance that was required to gain initial access?

Agentless EPS – An Alternative

Agentless endpoint security is network access control without the hassle of trying to maintain legions of agents on hosts that travel the country, often only connecting for minutes at a time. That's not to say that agentless EPS is as simple as plug-and-play, but nothing secure ever has been. Agentless EPS replaces agents on host machines by conducting detailed security scans of new assets and providing an objective compliance status from which the access control device can determine what level of access to grant.

Agentless EPS starts with an external policy validation server. Its job is to maintain an extensive database of policies that it can apply to new hosts based on the OS and installed software.

How Agentless EPS works:

1. When a new host, without an agent, requests access, the NAD sends a request to the ACS for compliance validation.
2. The ACS sends a request to the external policy validation server to scan the asset and return compliance data.
3. The scan engine scans the device and sends the information back to the external policy validation server where compliance policies are applied and an overall security posture token is assigned and sent back to the ACS.
4. Based on the posture token, the ACS tells the access device whether to permit, deny or limit the new host's access to the network.

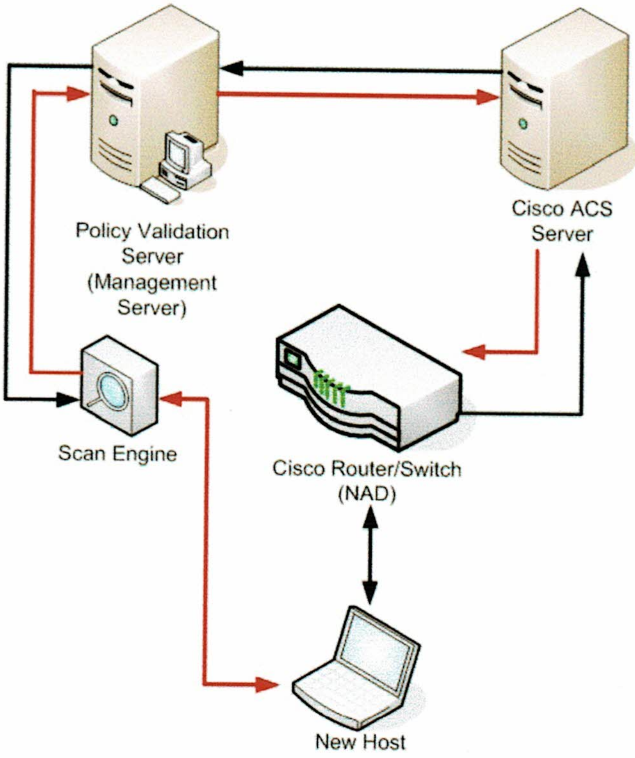
Policy Validation Server (Management Server)
Cisco ACS Server
Scan Engine
Cisco Router/Switch (NAD)
New Host

With agents, almost all of the intelligence is built into the agent, and the more information it must collect and analyze, the more resources it requires. With an agentless solution, the tasks are divided up between an external collection device (scan engine) and the management server, requiring almost nothing from the asset itself from a resource standpoint.

Another advantage of an agentless EPS is that it works for assets, such as laptops from contractors and service providers, that cannot have one of your agents on them. If devices from other organizations need a pre-admission scan before they are allowed to join your network, an agentless EPS solution is required. Further, once you are using an agentless EPS for this type of pre-admission scanning, the marginal cost of scanning your own assets before they join the network is minimal.

What to Look for in an Agentless Solution

Not having to worry about agent and software interactivity, resource usage, and how to guarantee agents are up-to-date makes the selection of an agentless solution a much easier sell; but selecting the right agentless solution still requires some due diligence.

The main things to examine when evaluating agentless EPS solutions include:

- The scan engine
- The compliance policy library
- Frequency of compliance evaluation

THE SCAN ENGINE

In the absence of an agent, the mechanism for collecting the information from an asset lies in the scan engine. The scan engine scans an asset, such as a laptop trying to join the network, and sends the information back to an external policy validation server where compliance policies are applied and an overall security posture token is assigned and sent back to the ACS. Some organizations are able to use one scan engine for an entire network, and others use several scan engines placed at different points around the network. For instance, agentless EPS wouldn't be secure if it had to raise and lower firewalls to allow a scanner to continuously scan remote segments of the WAN, so a distributed approach to scan engine placement would be required.

THE COMPLIANCE POLICY LIBRARY

Perhaps the most important piece of any security compliance solution (not just an agentless EPS solution) is the policy library. Basically the brains of the entire agentless compliance solution, the policy library is where an organization can define its requirements. Further, health tokens are generated based on whether assets are in compliance with the policies in the library.

A superior compliance policy library should consist of established, industry standard-based security policies, such as those developed by the National Institute of Security Standards (NIST), and written using an accepted language such as the Open Vulnerability and Assessment Language (OVAL) (XML-based) so that it can be easily expanded to fit the needs of the organization.

Your agentless EPS solution should be able to use this policy library to scan for a wide variety of configuration variables. Sure, you've got to do the standards, like patch level, hotfix level, and antivirus/anti-spam update level. But can you also check things like security vulnerabilities, security audit settings, and file integrity? It's important that the agentless solution you select can check and monitor any of the system configuration or operation variables that you consider important.