

Collaborative Endpoint Security (part I)

Iván Arce, Eduardo Arias
Core Security Technologies

Protecting endpoint systems such as desktop computers and servers is an important part of any reasonably well-thought security strategy for both enterprise networks and Internet user's home computers. The outbreak of devastating worms and email-borne viruses of the past years and the constantly increasing annoyance – and damage — brought in by SPAM and spyware echoed in the public ear the mantra that many security experts have been chanting (and perhaps mumbling and yelling) for more than a decade: "...defense in-depth, defense in-depth, defense in-depth..."

It now seems that the mantra has been heard and endpoint security is a serious concern for many security-conscious users. Multiple endpoint security solutions converged on the desktop from the perspective of feature set integration and security product or service offerings from a myriad of security vendors. Host-based intrusion prevention systems (H-IDS) seem to be the rising star of this pack, yet very few innovative ideas have seen the light of day in terms of how to deploy and operate them, as well as how to determine the associated value-model they propose.

In part I of this article we introduce endpoint security solution technologies and analyze some of the technical challenges they face in providing effective security to Internet users and organizations. A collaborative approach that

relies on cooperation between not only software components, but also between the users of endpoint security solutions is proposed as a plausible way to address these challenges.

Why endpoint security?

It is evident these days that the traditional perimeter defense approach to information security is helpless against a myriad of attacks in any but the most simplistic network setup. From a purely “theoretical” perspective the proposal for a strong defense system at the network perimeter – built mainly using firewalls and network devices such as routers and switches with certain security features – lacks, by definition, the visibility and depth to provide effective security in the rapid changing landscape of today’s software and network topologies. Undeniably, firewalls and other perimeter defenses provide some good security countermeasures to prevent attacks from passing through them when it is possible to inspect and sanitize inbound and outbound network traffic. The total network security advocate would argue that perfect (or quasi-perfect) network segmentation and enforcement of security policies at the perimeter could prevent most, if not all, security incidents. However, the basic founding premise for this security strategy remains the same as in the early ages of information security: A hard network shell and a relatively softer network core. The “security in-depth” school of thought would claim that this is “philosophically” insufficient to achieve a good security posture, as several layers of security mechanisms with decreasingly or increasingly stringent controls would provide a stronger overall defense against both external and internal attackers. The rationale is that, by combining several layers with various degrees of strength, the overall robustness, redundancy and effectiveness of the security infrastructure is increased and its deployment is pervasive to the entire network that it aims to protect (including end systems).

On the other hand, according to host-based security proponents, the required number of network security layers needed to deploy reasonable security adds up to be impractically complex, unmanageable or expensive (or all of the above). Even then, the host-based security proponents go, these multiple layers of network security would still fail when facing attack vectors invisible to outer network security layers, for instance when the internal network core is directly targeted, and therefore real security requires mechanisms deployed at the endpoint systems. Relatively recent security initiatives (e.g., Cisco’s NAC, http://www.cisco.com/warp/public/cc/so/neso/sqso/csdni_wp.htm or Microsoft’s NAP, <http://www.microsoft.com/windowsserver2003/techinfo/overview/napoverview.mspx>) seek to combine both worlds in a seamlessly integrated and interoperable manner. However, these initiatives are yet to prove effective, suitable and of real value in live, real-world scenarios. Meanwhile, endpoint and network security threats and solutions continue to evolve at a rapid pace and, as common sense would dictate, a good mix of host and network-based components are generally used in the security infrastructure of typical IT environments.

Threat prevention at the endpoint

That endpoint security software is a major portion of today's security infrastructure is demonstrated by the fact that anti-virus software is the most mature security technology and the most well-established vendors of information security dominate this market segment. Undoubtedly, anti-virus software packages are the most widely deployed security solution in IT environments across the board. Still, past decades have demonstrated that anti-virus software alone is not enough to cope with the emergence and evolution of new security threats.

Accompanying the explosive growth of computer networks – and perhaps somehow fostered by it –, users and organizations have increasingly turned to network-based security solutions in their search for better security controls that could complement, or even replace AV software.

This triggered the rise of the firewall and, later, of the network IDS as de facto, mandatory components in most security strategies.

As was the case with the previous endpoint-centric AV solutions cycle, pure network security plays proved insufficient to cope with security threats after the emergence of automated massive attacks, worms, directed attacks using exploit code that passes unnoticed through network devices, and with the various forms of malware targeted at endpoint systems and its users. The attention turned back to endpoint security solutions and consequently firewalls and Intrusion Detection or Prevention Systems adapted and moved into the desktop computer taking the form (or just the name) of Personal Firewalls, H-IDS/IPS or endpoint security policy enforcement solutions. The anti-virus software itself mutated and evolved to cope with many of those new threats.

Today, malware detection and removal, SPAM and pop-up blocking, protection against exploitation of software flaws (mainly against code-injection exploits) and application sand-boxing have all converged at the endpoint and a multitude of different – and sometimes conflicting – solutions are available from a similarly large number of possible providers. All of them face non-trivial technical and operational issues that they need to address in order to be successfully deployed and to provide effective defense at the endpoint. In the next section we will provide a laundry list of known issues that we consider that an endpoint security solution should address.

The visibility issue

Network based solutions can't detect or prevent what they can't see. For what they can see (i.e. network traffic data or traffic metadata), they lack accurate context to correlate observations to actual events on endpoint systems (i.e. why, when and how is the observed traffic generated and what generated it). The flip coin to this problem is that pure endpoint security solutions are not "network aware" and therefore lack context at the global or even local network level. A given endpoint solution can see what is going on at the endpoint system where it is running but lacks the network visibility to understand the aggregated effect of multiple endpoint systems with similar behavior. This imposes some fundamental limitations on what contextual information the endpoint security solution can act upon and the effectiveness of its security posture

The effectiveness issue

As the last line of defense against attacks, endpoint security solutions should (or, in a perfect world, must) be effective in detecting and preventing all types of attacks: those that are publicly known (associated to a publicly known vulnerability, exploit or “worse-practice”) and those that are yet unknown to the general public (attacks that exploit 0-day vulnerabilities or use new attacking techniques).

A flaw in the solution’s effectiveness can lead to direct security compromise in the case of false negatives (the case when a harmful event is considered harmless and not acted upon by the security mechanism), or to a disruptive and possible harmful reaction in the case of false positives (a harmless event that is considered harmful and acted upon by the security mechanism). An effective solution must constantly execute a delicate balancing act between the overzealous and over permissive extremes. Tilting towards one extreme renders the solution ineffective and self-defeating for its security purpose; tilting towards the other renders the solution unmanageable, untrustable and ultimately equally self-defeating.

All current endpoint security mechanisms make their stand somewhere in-between those two extremes by taking one or many architectural decisions.

The simplest approach is to use signature-based detection of harmful events. This is the traditional approach of AV software and its effectiveness depends on both the quality of the signatures and the pace at which new signatures for new known threats can be distributed. To cope with the unknown, AV software (as well as other endpoint security solutions) has adopted heuristics-driven technology to trigger security defenses. By using heuristics that define and detect known techniques and malicious behavior, endpoint security solutions can act upon known **attack patterns** that target components not necessarily known to be vulnerable.

Finally, the subsequent evolution of endpoint security technology led to almost purely behavioral models to define what is the normal security context within a protected system and act upon deviations from normality that surpass a given threshold, built either from preset settings or by learning what is ‘normal’ along its runtime. In this last case, the effectiveness of the solution derives from its ability to correctly model normal behavior and to clearly distinguish deviations. Behavioral solutions also rely on the assumption that attacks are events distinguishable from normal legitimate use of system resources. The behavioral approach offers the advantage of apparently coping with the unknown in a more effective manner and not relying on fast generation and distribution of configuration updates to prevent malware outbreaks.

The intrusiveness issue

Endpoint security solutions must also be easily deployable, manageable and should not conflict with normal use of the system by end users. The effective control of system resources needed to detect and prevent attacks often also requires using components that are quite intrusive at the OS level (i.e., kernel drivers or kernel modifications). On the other hand, less intrusive solutions may be circumvented with simple variations of attack techniques. Another delicate balance must be found between overly intrusive and trivial-to-evade technology to achieve security effectiveness and manageable deployment and operation.

The granularity issue

A great deal of granularity is needed to implement effective security control on system resources (this is specially the case for personal firewalls, HIPS and application sand boxing solutions that rely on configuration settings to control exactly what, when and how are the system resources accessed or used). In lieu of generic security mechanisms or purely behavioral-based solutions this greater granularity poses a difficult configuration problem: it is quite hard to know exactly which resources and what use of them is needed for all the possible legitimate uses of the endpoint system by all possible end users that have access to it. To cope with this problem, security providers:

1. Purposely decide to lose granularity in exchange for configuration and operational ease;
2. Build multiple abstraction layers to group granular permissions into objects or other abstract entities that are easier to understand and manage; or,
3. Transfer the problem of configuring highly granular permission to end users

Clearly, the first approach may introduce weaknesses due to the lost of security capabilities in the solution. The second approach maintains the architectural capabilities of the solution, but hidden weaknesses may be introduced if invalid or shortsighted assumptions are made for the abstraction layers and the grouping of granular permissions. The third option just transfers the issue to somebody else. Total control of the security solution is still possible through the painstaking process of configuring every single security capability available, but the end result is usually flawed due to the limited time and other resources that can be invested in such a project (also this often the case with the observed lack of precise security configuration for network firewalls and IDS).

The preparedness issue

Another very concerning issue for an endpoint security solution is that of its preparedness to cope with the threat of massively automated attacks in a proactive manner.

A hard-learned lesson of the past 5 years is that endpoint security solutions must be suitable to prevent the outbreak of fast propagating worms that exploit known security vulnerabilities in endpoint operating systems (such as CodeRed, Nimda, Blaster and Sasser) or that rely on malicious email-borne content that tricks the end user into enabling self-replication and propagation (LoveLetter, Bagle, MyDoom, NetSky, etc.).

To accomplish this goal, endpoint security solutions rely on one of two possible strategies:

- Signature matching or otherwise security configuration dependant solutions that rely on being able to quickly receive configuration updates that are general enough to detect or prevent exploitation of vulnerabilities for all – or as many as possible – variants of known exploitation code and malware. This strategy requires fast detection, analysis and response from the solution's provider and an equally fast and robust distribution channel.
- Behavioral and heuristic-based systems relying on the assumption that any new threat will be detected and prevented by the endpoint solution because, no matter what the new vulnerability or malware is, it will follow known exploitation techniques or behave in noticeably different manner from normal use of system resources.

Both strategies have strengths and weaknesses and have yet to prove effective in preventing global outbreaks of fast propagating attacks while keeping endpoint resources available for normal use. Behavioral and heuristics based solutions seem promising in this area, but their install base is still not sufficient? to prove effectiveness at a global scale.

Usability

Endpoint security solutions face a dilemma we hinted at in previous sections. A given solution must be not only granular and robust enough to detect and prevent attacks effectively, but also easily usable for a vast range of users with different degrees of technical skills and security-awareness. Complex configuration chores and the solution's internals must be hidden from end users, yet they must be good enough to cope with sophisticated attackers and the tools they use. The user interface of endpoint security solutions clearly epitomizes a debate (Security vs. Usability) that has been raging for decades in the various disciplines that deal with software engineering and information security.

Addition and distribution of value

In the rapidly changing information security landscape of our days, an endpoint security solution must be able to effectively cope with new attack techniques and future attack trends. For that purpose, a key feature is the ability for its users to improve and refine the capabilities of a deployed solution, which is equivalent to the ability to **add value** to the solution. Addition of value by users can take various forms but almost always implies the ability to incorporate new configuration settings (signatures, permissions, capabilities, etc.) in a way that makes sense within the user's IT environment (network topology, endpoint security policies and mission-critical applications and business processes). The suitability of a given solution for the addition of value by the end user is a variable with a significant influence in the effectiveness of the solution and its overall total cost of ownership (TCO). Ideally, an endpoint security solution should not only provide a way for end users to add value to it, but also provide the ability to re-use the added capabilities across a local (or better yet, global) installed base. The requirement for being able to re-use or **distribute added value** poses additional software engineering and information security challenges, since endpoint systems are usually the most diverse computing environments on any but the most simple networks: File system hierarchies, installed applications, files and directory names, registry hives and endpoint resource's usage practices vary largely even on local networks and are heavily reliant on end user idiosyncrasies. As of today the authors are not aware of any endpoint security solution that clearly addresses the addition and distribution of value by end users.

The rationale for a collaborative security approach

A collaborative approach to endpoint security may be suitable to address many of the issues described in the previous section.

The use of an endpoint security solution based on software agents that share context information about endpoint security state could provide the required network-level visibility of network devices.

Several proposals for collaborative agent architectures applied to information security have seen the light in academic circles during the last two years. These are

generally aimed at using collaborative agents to augment context visibility of network security appliances or to provide more efficient means of detection of global scale attacks. However, they are focused on **collaboration between software components** and not necessarily on **collaboration between security software users**.

Collaboration between users of security software can go beyond the technical advantages of peer or grid computing if a reasonable number of information security practitioners step up and actively collaborate to tackle the complex problem of implementing effective, usable and easily manageable endpoint security solutions. The success of several well-established security projects such as the Nessus vulnerability scanner and the Snort network intrusion detection system (rooted in the community-oriented development processes that distinguishes open source projects), are a good indication of the collaborative approach as a viable option. However, it should be noted that Nessus and Snort were not specifically designed to facilitate generation and distribution of security value in a collaborative manner, although both had a relatively good success at that.

Unfortunately, both projects are currently not within the scope of endpoint security solutions and they are not likely to face many of the challenges posed by the issues we described in the previous section of this document. Consequently we can't derive better-informed conclusions about the possible application of the existing community-oriented open source security projects to collaborative endpoint security.

Let's examine the possible advantages of a collaborative approach to endpoint security software and how those issues could be addressed.

Collaborative software components and users can leverage endpoint security technology to acquire a broader view of attack trends at the local or global network scope. Automatic sharing of security alerts, normal and abnormal behavior patterns of commonly used applications and configuration settings for endpoint security policies can provide improvements to the visibility and effectiveness of endpoint security solutions. Behavioral models can be enriched and improved with the addition of globally-seen behavior rather than just local execution patterns. Furthermore, with the proper tools, collaborative endpoint security users can improve the effectiveness of their solutions through the use of refereed or peer-reviewed endpoint security settings considered effective by one or more groups of users with specific skill sets, geographical or topological distributions and functional roles among the entire user base.

The granularity issue (a desired security feature for tight access control to resources in endpoint systems) can be approached with a divide-and-conquer strategy that distributes among collaborative end users the problem of defining very granular security permissions for the correct – and secure – execution of the endpoint's operating system and the overwhelming number of client applications in use today. Configuration tools that support and promote the seamless addition and distribution of value to endpoint security solution deployments must facilitate this distributed approach to security configuration.

In this fashion, end user experience and expertise with third-party software and applications used by specific user groups, vertical markets or networks of organizations can be leveraged to provide effective security settings to entire user base across geographical, topological or organizational boundaries.

Collaborative end users can achieve the optimal balance for intrusiveness, granularity and usability of endpoint security solutions in real world environments through the open discussion and improvement of security configuration settings. Voting, auctioning and “fair market valuation” techniques for endpoint security configuration settings can help determine what is the right balance for specific operational scenarios on local and global scales.

Finally, the preparedness issue can be addressed by a network of collaborative agents and users that provide the means for fast propagation or configuration settings in either reactive or proactive ways at the outbreak of global or local threats. For that purpose, a community-oriented website, instant messaging, P2P technology and traditional email and file transfers can be used.

Endpoint systems are increasingly vulnerable to security threats to the point that they’ve become the weakest link in the security posture of organizations and Internet users. Client-side exploits and *phishing* attacks, SPAM, virus, worms, *rootkits*, key loggers, distributed denial of service (DDoS) agents and other malware are a present and real security threat that point at the need for an effective endpoint security solution.

In part I of this article we’ve proposed a collaborative approach to address the endpoint security threat. In part II we will present a freely available endpoint security solution that implements some of the features and ideas behind the collaborative endpoint security approach proposed.

If you wish to learn more about the Core FORCE project at Corelabs visit us at <http://force.coresecurity.com>

Ivan Arce is CTO and co-founder of Core Security Technologies.

Contact him at [ivan.arce {#at#}coresecurity.com](mailto:ivan.arce@coresecurity.com)

Eduardo Arias is Head Engineer at Core Security Technologies.

Contact him at [eduardo.arias {#at#}coresecurity.com](mailto:eduardo.arias@coresecurity.com)