Intelligent Whitelisting

An Introduction to More Effective and Efficient Endpoint Security

The volume and sophistication of malware is skyrocketing, and traditional antivirus approaches are struggling to keep up. Historically, the approach to dealing with the growing quantity and complexity of malware has been to build a better antivirus "mousetrap," without any shift in the underlying management model for vetting change in endpoint environments. The result has been bloated anti-malware technology with ineffective protection and abysmal performance. This increases endpoint total cost of ownership due to increased strain on IT resources and reduced end-user productivity, which puts further pressure on already flat or reduced IT budgets.

It's time to rethink how we protect our endpoints.

The typical security professional tends to look at endpoint control as a choice between black and white: the blacklisting signature-based antivirus technologies that struggle with today's threats or the first-generation whitelisting technologies that tend to impede user productivity. Think again, though.

A new "intelligent" approach to application whitelisting uses both methods and adds an automated way to determine whether the stuff in between – the so-called graylist – should be trusted and allowed onto your network.

Intelligent whitelisting provides a unified workflow that brings signature-based and behavioral detection together with the power of whitelisting capabilities, and adds a "trust engine" that controls what changes are allowed. This streamlines and automates the process of adding trusted applications to the whitelist. Intelligent whitelisting automates important queries against applications such as "Do I know where this came from?" and "Are others using it?" by using data from other endpoint security applications such as patch management to dial in the level of control and security desired. Not only does it reduce malware infection rates dramatically without affecting productivity, it also allows you to reduce the TCO of maintaining endpoints.

The Need for a New Approach

The sad fact that endpoint TCO has gone up while security effectiveness has gone down has been noticed by executives who hold the purse strings. Even as malware proliferation escalates, the C-suite is asking IT to clamp down on security costs. Today, there are three clear drivers that cause IT professionals to pause and rethink their endpoint protection strategies:

I. Exponential rise in volume and sophistication of malware

Between 2007 and 2010, the number of threats security researchers found needing new AV signatures had risen from around 480,000 per month to about 1.8 million per month.1 (See graph below) Not only has the volume increased rapidly, but so has the sophistication of attacks. Much of the malware is now designed by financially motivated criminal syndicates that develop malicious code to bypass antivirus defenses and to target specific organizations.

Unique Malware Signatures

II. Limitations of traditional approach

On average, AV software detection rates of new malware upon initial discovery are just 19%, and a mere 62% after 30 days.2 Traditional antivirus came of age during a very different era, when malware variants remained limited. Vendors of signature-based antivirus have struggledto keep up with the exponential growth in malware, to no avail. Not only has the signatures glut degraded endpoint performance, but the typical antivirus vendor can no longer keep up with the surge of new and sophisticated variants. So, many slip through the cracks.

III. Budget constraints and increasing endpoint TCO

CIOs and CFOs are holding the line on IT spending in the near future. Gartner predicts that organizations will reduce the share of security spending by 3% to 6% of their overall IT budgets through 2011.3   That's bad news for organizations stuck in endpoint security status quo. The current AV model has made it more and more costly for IT departments to keep endpoints infection-free. In fact, the average organization now reports upwards of 50 malware incidents that impact productivity per month,4 leading to an increase in IT help desk, incident response, and remediation costs.

Application Whitelisting and Application Control: What's the Difference?

If you're more than just a little confused about the difference between application whitelisting and application control, have no fear. The truth is that the two terms are separated by about as much difference as to-MAY-to and to-MAH-to. These two terms both describe the same process. We've just chosen to consistently use the term "application whitelisting" because "application control" has also taken on other connotations as application firewall vendors have searched for an appropriate term to describe their worthwhile – but different – product base.

The Shift to Intelligent Whitelisting

In its purest sense, application whitelisting turns the traditional antivirus approach on its head. Instead of poking at a suspect piece of software and looking to see if it's bad in every which way possible, application whitelisting asks the fundamental question, "Do I have reason to trust this code?" At its heart, a whitelist solution is seeking to confirm an application as a valid piece of software. Until it does so, that application cannot run on the endpoint. In the most simple of deployments during whitelisting's early days, all executions of application and code were limited to that verified list of known good code, which was an extremely solid way to keep malware off mission-critical servers. It then evolved as a security layer for "locked down" endpoints, such as retail environments with point-of-sale (POS) systems or call center environments.

However, most workers today operate in a much more complex and dynamic environment. More applications are downloaded and used to perform job duties, including open source tools, web applications, home-brew code, and commercial programs—all of which change endpoint configurations and make them unique to each worker. Then add into this mix remote and mobile workers and the growing push to extend applications into the cloud. It becomes clear that whitelisting policies must be more flexible—both for IT and for end-users. In order to facilitate this kind of environment, whitelisting has to be able to offer enforcement that is adaptable enough to allow workers to safely leverage new tools that improve productivity.

Fortunately, forward-looking vendors have taken these issues into account and have developed smarter whitelisting solutions that offer better security, yet remain flexible enough for dynamic environments. Rather than constantly

managing a centralized whitelist before changes are allowed, intelligent whitelist users define a set of automated trust rules that are fine-tuned to their risk appetite and control tolerance. This eliminates the need for constant intervention by IT, by automating the verification of good software using common indicators such as the reputation of the software publisher or the reputation of the tool implementing an update or a new piece of software.

Managing Trust

Trust. It seems like a simple concept. Either you trust something to run on your endpoints, or you don't. But in reality, we know it's a lot more complicated. For instance, suppose you find a P2P application running in your environment. The files aren't corrupt, and it's a widely used program. But do you want it on your network? If you work at an organization with highly confidential information, then P2P applications probably aren't appropriate. On the other hand, if you work at an advertising agency that zips files back and forth with regularity, perhaps it's not such a bad thing. In the end, you should be in charge of that "security dial" so you can decide, instead of being limited by your technology.

In order to ensure improved endpoint protection, our intelligent application whitelisting solution should include the following capabilities:

Trust engine

Your intelligent whitelist solution should validate endpoint changes based on trust rules your organization establishes, and automatically update the whitelist accordingly. These trust rules should be flexible enough to allow you to validate based on:

- the publisher of the software, using digital certificates and other metadata;

- the updater that introduces new or updated software;

- the path or centralized location, to ensure that in-house developed or unsigned executables that change frequently aren't blocked; and

- local authorization for specific trusted users with a lot of unanticipated change needs.

Snapshots

A snapshot capability should allow you to create a local whitelist of all executables. By creating a local whitelist, you can prevent any further unwanted changes to the endpoint environment and eliminate the need for a "perfect" global gold image that you shoehorn your entire organization into. A snapshot capability also lets you speed up whitelist deployment significantly and roll up unique whitelists to the global level for central visibility, grouping, and policy assessment.

Control over "local admin" users

In many organizations like yours, end users are granted local administrator privileges to ensure they have the flexibility to install and run application updates in order to get their jobs done. This approach to end-user management has led to chaos, resulting in a complete lack of control over endpoint configurations. This leaves systems much more vulnerable to exploits.

In contrast, intelligent whitelisting allows your users to maintain their local admin roles, but places limits on the kind of changes they can make and how much access they have to local system consoles that affect configuration changes. The end result is a more productive end user, while you obtain greater visibility and control of your desired endpoint security configuration and posture.

Ability to fit into an overall endpoint management workflow

Application whitelisting is only intelligent if it's easily layered into an overall framework that includes a spectrum of other endpoint security and management tools. By consolidating information that has traditionally been siloed off into these different types of tools and by also looking at any provenance and prevalence information, you can significantly enhance endpoint protection. While antivirus may be losing its effectiveness as a stand-alone solution, it is still a valuable tool when the information that is produced and stored within it is paired with application whitelisting. Similarly, if the information stored within a whitelisting tool can be integrated smoothly with patch management and trusted change policies, you can improve your organization's security posture and also reduce TCO for endpoints.

Creating Trust Based Policies:

An intelligent whitelisting solution will help you answer these questions:

- Is this known bad?

- Is this known good?

- Is this unwanted?

- Is this authorized?

- Is this properly licensed?

- Do I trust the vendor?

- Do I trust the program that introduced it?

- Do I trust where it came from?

- Do I trust this user to install it?

The Issue Isn't Black or White

Deciding to use application whitelisting is no longer an 'either-or' decision. The choice is no longer between antivirus and whitelisting. As Nigel Stanley, analyst at Bloor Research puts it,"I think the problem with whitelisting and blacklisting is that, superficially, it is too black and white! Of course there is a range of code out there which can easily be deemed to be nasty and is easy to blacklist. Similarly, there is code which is easier to whitelist—think downloads from major software suppliers. That said, I have known what appears to be "goodware" downloaded from a trusted vendor which then promptly screws up an IT estate due to application compatibility problems. The merger of whitelisting with blacklisting is probably inevitable, with greylisted code sitting in the middle, maybe subjected to some heuristic analysis."[5]

Similarly, organizations no longer need to decide whether or not to use whitelisting based on how static or dynamic the environment is. Instead, the decision is about what policies will be used that balance flexibility and security. In very secure and static environments, you may want

to use a pure whitelist policy. In dynamic endpoint environments, you may trust the user but need to ensure that proposed changes are both trusted and authorized. In

the corporate network, perhaps new code that cannot be identified on a whitelist must be introduced by your systems management tools.

Either way, any smart organization should strive for a defense-in-depth approach that controls application deployment in conjunction with malware cleanup and automated patch management. This ensures that machines are configured securely, and that whitelists are properly and constantly updated. In the end, this blended trust-centric approach is both more flexible and more secure than current endpoint protection approaches. With this intelligent whitelisting approach, you'll witness:

- Improved security

- Reduced costs

- Better control over endpoints

- Improved productivity

This trusted change approach to endpoint protection lets you balance usability with security, increasing end-user productivity without adding IT administrative burden. And in today's IT environment, that's a good thing.

www.ivanti.com

1.800.982.2130

sales@ivanti.com

1. Extrapolated from: McAfee Labs, McAfee Threats-Report: Third Quarter 2010, November 2010

2. Cyveillance, Malware Detection Rates for Leading AV Solutions, August 2010

3. Gartner, Vic Wheatman, Research Director, June 2010

4. Ponemon Institute, State of Endpoint Risk 2011, November 2010

5 Nigel Stanley in Lumension blog, Winning the Malware Battle: The Move Towards Whitelisting, December 2009